

THE WORLD IN MIND

BY

NICK TREANOR

B.A. (HONS), QUEEN'S UNIVERSITY, 1998

A DISSERTATION SUBMITTED IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY
IN THE PROGRAM IN PHILOSOPHY AT BROWN UNIVERSITY

PROVIDENCE, RI
MAY 2009

© Copyright 2009 by Nick Treanor

This dissertation by Nick Treanor is accepted in its present form
by the Department of Philosophy as satisfying the
dissertation requirement for the degree of Doctor of Philosophy.

Date _____
Jaegwon Kim, Advisor

Recommended to the Graduate Council

Date _____
Christopher Hill, Reader

Date _____
Ernest Sosa, Reader

Approved by the Graduate Council

Date _____
Sheila Bonde, Dean of the Graduate School

Nick Treanor was born August 28, 1972 in St. Catharines, Ontario. He attended Queen's University in Kingston, graduating with a first-class degree in philosophy in 1998. He has published “The Cogito and the Metaphysics of Mind” (*Philosophical Studies*, August 2006) and is currently working on papers on consciousness and on how the world is represented within thought. His “The Measure of Knowledge” was selected by the 2008 Young Philosophers Lecture Series. He has taught at Trent University and the University of Toronto, and in September 2008 will take up a five-year college lectureship at the University of Cambridge.

Acknowledgments

In my first or second year at Brown I took a graduate seminar course with Jaegwon Kim in which we read, among other things, Thomas Nagel's *The View from Nowhere*. Such seminars typically involved Kim working a small way through a long list of notes and questions he had typed up as he read the work in question. One day, Kim spent a lot of time asking “What is a point of view?”. His tone suggested genuine puzzlement. I recall thinking, “What is this guy's problem? Has he really never heard the expression 'a point of view' before? How is that possible?”. Only much later did I come to recognize that I, too, had no idea what a point of view is.

Kim's seminars were distinguished by his remarkable ability to recognize what we don't understand. But Kim also set a stunning example of philosophical integrity, for he was always patient with the problem at hand, charitable with texts, and completely devoid of the vanity that threatens intellectual sincerity. In many places within the present work I fall far short of this ideal, but I am indebted to Kim for helping me see what philosophy ought to be. I am also grateful to Kim for giving me the freedom to pursue the topic of this dissertation, for his encouragement and confidence, and for penetrating comments on early drafts.

Chris Hill gave shape and direction to the work by his early insistence that I am crazy to think that a subject has a global representation of the world. This helped me see that there is a genuine problem concerning the many/one structure of belief. Hill also provided very helpful comments on the penultimate draft of the dissertation, raising issues and asking questions that I will think about for a long time to come.

I have also benefited from many seminars with Ernie Sosa. It was in one such seminar that I came alive to the problem of understanding the goal of inquiry, which is an early motivation for the current work. Sosa, too, helped me understand that real depth in philosophy requires tenacity.

The philosophy department at Queen's University at Kingston, where I studied as an undergraduate, gave rough form to my philosophical sentiments, and I am grateful for the patience and grace its members showed toward me. Here I should mention in particular Steve Leighton, who knows why.

Trent University and the University of Toronto provided me with academic employment as I finished the dissertation, on generous and respectful terms.

I learned much about philosophy from my dear friend Derek Ettinger, and I am gratefully especially for his patient support throughout my time at Brown. It was in the course of a conversation with Derek in the spring of 2005 that I first came up with the idea that beliefs are bound together by diversity judgments.

This idea was refined after a disastrous job interview with York University in Toronto, in which I unwisely tried out ideas for the first time. Warm thanks to Henry Jackman and Jagdish Hattiangadi for showing me the problem.

But my two greatest debts are: to my parents, to whom I owe so much more than they realize, and to Lisa, the most gifted researcher I know.

For mom and papa

When we first begin to *believe* anything, what we believe is not a single proposition, it is a whole system of propositions. (Light dawns gradually over the whole.)

– Ludwig Wittgenstein, *On Certainty*

It is a rather curious fact in philosophy that the data which are undeniable to start with are always rather vague and ambiguous. You can, for instance, say: ‘there are a number of people in this room at this moment.’ That is obviously in some sense undeniable. But when you come to try and define what this room is, and what it is for a person to be in a room, and how you are going to distinguish one person from another, and so forth, you find that what you have said is most fearfully vague and that you really do not know what you meant. That is a rather singular fact, that everything you are really sure of right off is something that you do not know the meaning of, and the moment you get a precise statement you will not be sure whether it is true or false, at least right off. The process of sound philosophizing, to my mind, consists mainly in passing from these obvious, vague, ambiguous things, that we feel quite sure of, to something precise, clear, definite, which by reflection and analysis we find is involved in the vague thing that we started with, and is, so to speak, the real truth of which that vague thing is a sort of shadow.

– Bertrand Russell, *The Philosophy of Logical Atomism*

TABLE OF CONTENTS

Chapter 1: The Unity of Belief: An Introduction	1
Chapter 2: Some Initial Objections	33
Chapter 3: Davidson's Holism	62
Chapter 4: Lewis and Stalnaker	102
Chapter 5: Unity and Conjunction	132
Chapter 6: Unity and Diversity	180
Bibliography	203

THE UNITY OF BELIEF: AN INTRODUCTION

A philosophical account of judgments lies, however, in their having a further significance than merely being right or wrong: they build up our picture of the world.

– Michael Dummett, “Truth and Meaning”

The question “What is the goal of inquiry?” can be given either a descriptive or a normative reading. The descriptive reading is one on which the question is asking what goal cognizers in fact have. Naturally the descriptive question can be refined in various ways: Is there something all cognizers have as their goal in inquiry? What does a cognizer typically have as her goal in inquiry? What does this particular cognizer have as her goal on this particular occasion? But the question can also be given a normative reading, on which what is at issue is what inquiry, as an activity, takes as its principal or defining aim. What does inquiry, in its nature as inquiry, aim at? This is the more philosophically interesting question and the one philosophers typically have in mind when they discuss the goal of inquiry.

My purpose in this chapter is not to engage significantly with the normative question of what the aim of inquiry is. But I want to begin by looking briefly at the issue, as a way of drawing out what I take to be an important ambiguity, or duality, in how we think about our cognitive relationship with the world. We can think of the mind-world relation as a relation between parts of minds (individual beliefs) and specific features of the world (whatever state of affairs, in the broadest sense, an individual belief represents as being the case). Or we can think of the mind-world relation as a relation between whole belief systems and the world as a whole. My intent is to clarify and understand the connection between these two kinds of

mind-world relation. This dissertation focuses on just one aspect of this issue – a problem in the philosophy of mind concerning the many/one, or part/whole, structure of belief. It is a project within the philosophy of mind, but the story begins, as it were, with epistemology.

The Goal of Inquiry

Consider the following remark by Laurence Bonjour, which appears in the opening pages of

The Structure of Empirical Knowledge:

what makes us cognitive beings at all is our capacity for belief, and the goal of our distinctively cognitive endeavors is truth; we want our beliefs to correctly and accurately depict the world. (1985, p. 7)

Bonjour's remark is a platitude, the kind of general claim with which one often begins a philosophical investigation, before the serious work has started. I cite Bonjour not to focus on him in particular but because his claim that we want our beliefs to accurately depict the world is a familiar one: we have all heard this claim, or something similar, before, and most of us endorse it. Yet for all its familiarity, the claim is ambiguous between two quite different interpretations.

How is Bonjour's claim ambiguous? Consider two distinct things someone using that expression could have in mind:

1. We want each of our beliefs to correctly depict some feature of the world.
2. We want our beliefs, as a whole, to faithfully represent the world.

These two different interpretations are themselves not yet clear, for we still need to know what "correctly" means, what "faithfully represent" means, what "feature" and "as a whole" mean, what "depict" means, what "world" means, and so on. Nonetheless, the two interpretations mark a general difference, two different paths, the distinctness of which would survive a clarifying analysis of the individual terms. The first emphasizes the mind-

world relation as a part-to-part relation (individual beliefs to specific features of the world), where what is desired is that each belief fit, or correspond to, some particular state of affairs. The second interpretation, in contrast, emphasizes the mind-world relation as a whole-to-whole relation (whole belief systems to the world as a whole). Here what is desired is that one's belief system fit, or correspond to, the world conceived of as one entity.

That the two interpretations are genuinely distinct becomes clearer when we notice that a cognizer could satisfy the first desire without satisfying the second. Consider a very simple cognizer, say a frog, whose representational capacity and accomplishments are presumably severely limited. A frog may well depict the world accurately in the first sense, for it may be that each of its representations (of flies, its immediate environs, various states of its body, such as hunger, and so on) is veridical.¹ Yet it can hardly be thought to depict the world accurately in the second sense, for so much lies beyond its ken. Each representational state it has corresponds to, and is a good fit to, some feature of the world; but its representational system as a whole falls far short of providing a faithful representation of the world as a whole.²

The same point holds for human cognizers – the mere fact that the beliefs we have are true does not itself ensure that our beliefs, as a whole, faithfully represent the world as a

¹ It is not clear that a frog's representations could be *true*, in the strict philosophical sense, for it is not clear that they are propositional. It is standard to assume, however, that a frog's representations can be, and can fail to be, veridical, where this is a more general concept of correctness.

² Rescher provides a good characterization of the distinction, although his terminology does not match my own: "Understanding aims not only at the truth but at the whole truth....It hinges on how matters stand "with everything taken into account". Consider an analogy. I am going to buy a car, and want to know if it is economical to operate. It might be very economical in part of fuel usage but not (say) in terms of its demand for lubricants or in point of frequency-of-repair considerations. If I want to know if it is "economical on the whole -- *with everything taken into account*," then I must combine all these aspects into a single, comprehensive, over-all evaluation.

This achievement of adequacy in understanding is a matter of combining points of view and synthesizing them into a unified whole. It is to be achieved through reproducing the aspectival complexity of the object on the cognitive side through a comprehensive perspectival characterization. *Adequacy lies in wholeness*: it is only achievable through comprehensive synthesis of diverse aspects." (1979, 24, italics in original)

whole.³ The point holds also of our representation of objects within the world. If I were to say, for instance, that I want my beliefs to accurately represent some frog with which I am acquainted, I might mean that I want my belief system to be such that every belief I have about that frog is true – that it is green and alive, that it developed from a tadpole, that it has never performed differential calculus, and so on. Or I could mean that I want my belief system to be such that it gets the frog right, where this means more than that it contains at least one true belief about the frog and no false beliefs about the frog. It is no easy task to say what exactly it is to faithfully represent the frog in this second sense, but some part of the idea seems to be that, in general, the more I learn about the frog, the closer I come to satisfying the desire to faithfully represent the frog in this second sense.⁴

The dual sense of “accurate representation” applies not merely to our representation of the world as a whole, and of the representations we have of objects, animate or otherwise, within the world, but also to our representation of such things as animal husbandry, plumbing, French cuisine, the history of the Peloponnesian war, the works of Gerard Manley Hopkins, the prevailing theories of the physical sciences, or the truths of metaphysics. I have, to choose one of these as an example, various beliefs about Hopkins’s oeuvre, where these beliefs, collectively, depict Hopkin’s oeuvre as being some particular

³ Cf. Rescher 2003, p. 48: “Ignorance, lack of information, cognitive disconnection from the world’s course of things – in short, errors of omission – are also negativities of substantial proportions. This too is something we must work into our reckoning.”

⁴ I hedge here and include “in general” because it is not obvious that *everything* I learn about the frog improves my overall representation of it. Suppose I know, for instance, that frog F, the one in Providence I’m thinking of, is green, and then later, while in Vancouver, acquire the concept “Stanley Park”. I then learn (or come to believe) the proposition: Frog F does not live in Stanley Park. This is a fact about the frog, at least in some sense, but it is not obvious that it improves my representation of the frog. The inclusion of ‘in general’ is meant to limit the scope of the claim that when I learn more about the frog, my representation of it improves: although there may be exceptions to this as an absolute claim, it does seem clear and uncontroversial that my representation of the frog improves when I learn it used to be a tadpole, it can swim and move about on land, it weighs x grams, etc. (This point suggests that, eventually, we will need a principled account of the difference between facts about x the knowing of which improves one’s overall representation of x, and facts about x the knowing of which does not improve one’s overall representation of x. This may relate in interesting ways to the intrinsic/non-intrinsic distinction, but the issue is complicated and cannot be discussed here.) See “The Measure of Knowledge” (Treanor, in progress) for a discussion of related issues.

way. That representation could clearly be improved – not because it contains false beliefs, but because it is incomplete. I am familiar with some of Hopkins’s works, but not all of them, and not with every property of those works of his with which I am familiar. Insofar as his oeuvre is the intended object of my representation, the representation falls short of something with which I could rest satisfied were I to have as a desire the desire to be such that my beliefs faithfully represent the works of Hopkins as a whole.

To make this a little clearer, suppose I represent Hopkin’s oeuvre as consisting of “The Windhover”, “Pied Beauty”, and “God’s Grandeur”, and I have no beliefs about whether he wrote anything else. I believe that he wrote these three poems, I have beliefs about what each poem is about, I have beliefs about ways in which the poems are similar thematically, stylistically, and so on. Given that Hopkins did write these three poems, and supposing each of these other beliefs about the properties of these poems is true, all the beliefs I have about his oeuvre are true. Yet, given that many other properties of these poems escape my ken, and that Hopkins wrote a great many more poems, it is not the case that my beliefs about Hopkins’s oeuvre do a very good job of representing it. They certainly do not do nearly as good a job of representing Hopkins’s oeuvre as they do in the possible world in which those are the only three poems he wrote. In the actual world, my beliefs represent just a small part of Hopkins’s oeuvre, whereas in the possible world just described they represent much more of it.

Let us return to the broader issue, on which Bonjour’s familiar remark is ambiguous between the expression of the desire to have a belief system such that every belief it contains is true, and the expression of the desire to have a belief system such that it faithfully represents the world as a whole. One might happily grant that there is a genuine distinction here, that having one’s beliefs be true is not the same as having true beliefs that, as a

collective, do a good job of depicting the world, yet object that this second notion is not properly understood as a species of accuracy. That is, one might grant that inquiry has a two-part aim – the production of true beliefs, and the production of beliefs that, collectively, do a better job of representing the world – yet deny that this involves two different senses of accurate representation. A frog, this objection goes, does not do as good a job, qua inquirer, as we do, but this is not because we depict the world more accurately than does the frog but because we depict more of the world accurately than does the frog.

There are two things to say in response to this objection. First, accuracy is a veridical notion, and the heart of veridicality is *fit* or *faithful representation*. My belief system does not merely contain more accurate items (more true beliefs) – it itself is a better fit to the world than is the belief system of a frog. It is not merely that it gets more of the world right, it does a better job of getting the world as a whole right, where these notions are at least conceptually distinct. But I do not expect all readers to share these intuitions, and I recognize that they would need to be explored at length. Hence the second response: it doesn't really matter for present purpose whether we use the word 'accurate' in the way I wish to use it, as appropriate to characterize not merely the relation between individual beliefs and features of the world but also the relation between one's entire belief system and the entire world. The important point is that we can discern two distinct epistemic desires, regardless of what labels we use once we have distinguished them. A person with the first desire takes as her aim being such that each of her beliefs is true. A person with the second desire takes as her aim being such that her belief system, as a whole, is a good fit to the world, or represents the world faithfully, or more faithfully. Correspondingly, we can discern two different ways of thinking of the mind-world relation, regardless of the names we give each way. We can think of the mind-world relation as a relation between individual beliefs

and individual states of affairs. Or we can think of the mind-world relation as a relation between one's entire belief system and the world as a whole.

The first epistemic desire, and the corresponding conception of the mind-world relation as a relation between individual doxastic states and individual states of affairs, are familiar enough, for the issues they give rise to dominate the philosophical agenda. They principally divide into three. There are epistemological issues: how can one know or be assured that any given belief accurately depicts some feature of the world, and what makes some beliefs that accurately depict some feature of the world epistemically better than other beliefs that accurately depict that feature of the world? These are the traditional problems in epistemology of defeating the skeptic and of understanding the difference between mere true belief and knowledge. There are metaphysical issues: what is it for a belief to accurately depict some feature of the world? This is the traditional problem within metaphysics of understanding what it is for a belief to be true, or of understanding what truth is. And there are problems in the philosophy of mind: in virtue of what does a belief represent what it does; what makes it the case that my belief that p is a belief that p rather than that q , where p and q are different? This is the traditional problem in the philosophy of mind (particularly prominent today, but dating back at least to the medieval philosophers) of understanding what grounds or determines the content of a mental representation. (As Wittgenstein put a similar question, "What makes my image of him into an image of *him*?" (1953, p. 177)) These issues are amongst the most central in the philosophical tradition, and are those upon which so many other problems rest. They have deservedly been the subject of much philosophical interest over many centuries.

But what if – in addition to or as a consequence of our individual beliefs' depicting (accurately or otherwise) the world in the first sense – our beliefs, as a collective system,

depict the world in the second sense? That is, what if a representational relationship holds not merely between each individual belief and that feature of the world it is about, but also between the totality of one's beliefs, as a system, and the world as a whole? My view is that a belief system's so doing would raise interesting, related issues in epistemology, metaphysics, and the philosophy of mind – issues that parallel, but are distinct from, the familiar issues just mentioned. If our minds (or belief systems) depict the world in this second way, I believe, then rich new problems emerge as worthy of attention.

This is not a thesis I can defend in this dissertation, nor am I able to examine any of these 'new problems' in detail.⁵ My aim is more modest, and more basic. I will focus on exploring and attempting to make sense of the idea that a subject's beliefs comprise, or compose, something like an overall representation of the world. If there is a representational relationship between a subject's beliefs, as a system, and the world as a whole, this would seem to entail that a subject's beliefs, as a system, are or have as their content *a* representation. That is, let us grant that each belief represents some state of affairs as being the case; what is at issue is whether, and if so how, the totality of one's beliefs represents some state of affairs as being the case. At first glance the question may seem altogether puzzling. "How could the totality fail to represent a state of affairs as being the case if each of its members represents some state of affairs as being the case?", one might be tempted to ask. A fuller answer to this question will come later, but for now let me respond as follows. What is at issue is not whether the totality represents states of affairs, but whether the totality represents *a* state of affairs. By way of comparison: Suppose every Rembrandt painting

⁵ If a representational relationship holds between the totality of one's beliefs and the world, this bears on an issue of some interest in epistemology, that of whether inquiry aims not at the accretion of true beliefs, but at understanding. This latter notion is adumbrated as some kind of global state that involves seeing how things hang together, where this is typically understood as having some kind of internalist character. In my view it is more promising to go to the metaphysics instead: determine what it is for a belief system to better fit the world, and take that to be the normative goal of inquiry. Again, these issues are explored in "The Measure of Knowledge".

represents a state of affairs as being the case. Let us grant that, or set aside whether, the totality of Rembrandt paintings represents those very same states of affairs. The question to consider is: “Does the totality of Rembrandt paintings represent *a* state of affairs as being the case?”, or “Is there some state of affairs such that the totality of Rembrandt paintings represents it as being the case?”. Intuitively, the answer is no. If this is right, then the fact that each of a member of a collection represents a state of affairs as being the case does not entail that the collection represents a state of affairs as being the case.⁶ If, therefore, the totality of a cognizer’s beliefs collectively represent a state of affairs as being the case, an explanation or account of this is needed.

One might continue to insist there is no puzzle here, perhaps by saying that the relevant difference between the totality of Rembrandt paintings and the totality of a subject’s beliefs is that the beliefs belong to one cognizer. As I said, I will discuss these issues in more detail later, but I want to note at this initial stage, in reply to this kind of worry, that we are already familiar with two problems that parallel in interesting ways the problem I take to exist with regard to belief: the unity of consciousness and the binding problem. Considering these briefly may help to show what I take the problem concerning belief to be, and why the solution is not quite as simple as the imaginary opponent supposes.

As the unity of consciousness is usually framed, a subject has at a time many distinct experiences (of various colors, sounds, textures, and so on), but these experiences are deeply unified, in the sense that they are aspects or components of one overall state of consciousness. The binding problem is the distinct but related problem of understanding how a subject’s experience of objects within the world is built up from her experience of

⁶ Compare William James’s famous remark (in a different context): “Take a sentence of a dozen words, and take twelve men and tell to each one word. Then stand the men in a row or jam them in a bunch, and let each think of his word as intently as he will; nowhere will there be a consciousness of the whole sentence” (1950/1890, p. 160).

various aspects or components of those objects. Suppose a basket of apples, pears and peaches is sitting on the table in front of you. You have, it would seem, various distinct experiences (of the apples, of the pears, and of the peaches), where these experiences are aspects, in some sense, of the experience you have of the whole basket of fruit. The many individual experiences are, it is supposed, unified into one overall conscious state. Moreover, your experience of each of the apples, pears, and peaches, and of the basket itself, is, at least naively, built up from your experience of its various properties (its color, shape, location, etc.). Distinct cognitive mechanisms process perceptual information about these properties, and somehow this information, to speak somewhat loosely, is bound together into an overall conscious experience of the object. Both the unity of consciousness and the binding problem appear to involve a kind of compositional relationship between individual experiences or representational states and overall conscious experiences or representational states, and it is a real puzzle to understand the nature of this relationship.

The parallels with belief are striking. You have distinct conscious experiences of the various fruits within the fruit basket, and these conscious experiences seem to be aspects or parts of an overall conscious experience of the entire fruit basket (indeed, of your entire conscious state at the time). And your perceptual representation of each piece of fruit appears to be built up from your perceptual representations of various properties of that piece of fruit (its color, shape, location, etc.). Now compare the *beliefs* you have about apples, or pears, or peaches. You believe, I suppose, that apples grow on trees, that they are edible, that they have seeds, that they are usually red or green, that they are high in fibre, that they are widely cultivated. Each of these is an individual belief, but each is also, to put the point crudely, part of your conception of apples. Just as many distinct conscious experiences at a time appear to be tied together in the sense that they are aspects of a single encompassing

state of consciousness, many distinct beliefs at a time appear to be tied together in the sense that they are all aspects of a single, overall representation – of the world, or of an object or domain within it. It is no solution to the problem of the unity of consciousness to simply declare that a subject's many conscious experiences are unified into one overall conscious experience by virtue of being that subject's experiences. And it is no solution to the problem of the unity of belief to declare that a subject's many distinct beliefs are unified into one overall representation by virtue of being that subject's beliefs. (If anything, I would think the order of explanation runs the other way: A subject does not *give* unity to conscious experience and beliefs, a subject *falls out from* such unity.)

The chapter began with a discussion of the goal of inquiry, and noted that there is a distinction between having only true beliefs and having a belief system that faithfully represents the world. The former emphasizes a part-to-part mind-world relation (belief to individual state of affairs), while the latter emphasizes a whole-to-whole mind-world relation (entire belief system to world as a whole). I indicated that I thought significant philosophical issues arise when we take seriously the idea of the mind-world relation as a relation between an entire belief system and the entire world. I then narrowed the scope to a single topic – that of whether a subject has an overall representation in some sense composed of her many individual beliefs – and tried to give some initial character to the problem I will focus on in the chapters to follow. The balance of this chapter is devoted to three tasks:

- (i) Illustrating, via a brief survey of philosophical literature, how pervasive the concept of a single, overall representation of the world is within contemporary analytic philosophy. The notion is broadly used, by philosophers in different fields and with quite different commitments, although there has been little serious effort to understand it.
- (ii) Outlining, in somewhat more detail, the intuitive or pre-theoretical concept of what I call 'a global representation of the world'.

- (iii) Introducing two further concepts — ‘overall global representation’ and ‘global representation of x ’ – and distinguishing a number of separate issues and problems.

The first task is designed to motivate the project of the dissertation by showing that the concept I focus on is very much in use within philosophy, despite being not well understood. The second task is designed to establish in somewhat more detail what is meant by ‘global representation of the world’, which is the largely intuitive or pre-theoretical notion I hope to examine and better understand. The third task is designed to clarify and refine the focus of the dissertation.

The Pervasiveness of the Concept of a Single, Overall Representation

The concept of a single, overall representation of the world is already, I believe, part of our conceptual repertoire. In this section, I demonstrate how pervasive the concept is within contemporary analytic philosophy. The point of the brief survey is twofold. First, by showing how the concept is used, I can make somewhat clearer what I have in mind when I speak of a global representation of the world. Second, I hope to motivate the project of the dissertation by showing readers who are inclined to doubt the intelligibility of the concept of a global representation of the world, or doubt that a subject has any such thing, that I am not alone in thinking that a subject has a global representation of the world. What follows is not by any means an exhaustive catalogue of how the concept has been employed in recent philosophical work, but rather a sampling intended to convey something about the range and general character of the concept’s deployment. In some cases, the concept lies at the heart of a substantive philosophical position, such that it is not obvious the position would survive if it turns out that, after a full examination, the concept deserves to be abandoned. In other cases, it may be merely a *façon de parler*, one which could easily be eliminated if the

need arose. I will not attempt to explore which philosophers or philosophical positions are deeply committed to the intelligibility of the concept. There are at least many cases in which it is not obvious that the metaphor could be abandoned in favor of more familiar concepts, rather than clarified through analysis, without significant loss.

The concept of global representation of the world is employed by philosophers in distinct subdisciplines, and although it may not always be used in exactly the same sense, there is a common core of shared understanding concerning the concept. That these uses may not all be identical is not surprising, for the concept has not been, so far as I am aware, subject to serious examination. If the concept as used by one philosopher is not exactly as that used by another philosopher, this is perhaps because it remains vague enough to be used in distinct ways.

One place the concept is commonly found is in discussion of massive error, of the kind many believe would obtain if we were all brains in vats. Laurence Bonjour, for instance, says that if we were brains in vats and were fed beliefs and sensations which depict a wholly unreal world, “our picture of the world would be totally or almost totally false” (unpublished, section 2).⁷ Donald Davidson has famously set himself against this view, saying that “our view of the world is, in its plainest features, largely correct....What cannot be the case is that our general picture of the world and our place in it is mistaken, for it is this picture which informs the rest of our beliefs, whether they be true or false, and makes them intelligible” (1991, p. 160). What is of special interest here is that Bonjour and Davidson not only claim that we do in fact have a ‘picture’ or ‘view’ of the world, they claim that this representation is subject to something like veridical appraisal – it could be “almost totally false” or “largely correct”. Elsewhere Davidson is even more explicit that he takes the

⁷ Elsewhere Bonjour says that what he is calling ‘a picture of the world’ can be “in some important way inaccurate, incomplete, or distorted” (1985, p. 172).

global appraisal to be veridical. In a discussion of Nagel's reaction to his arguments, for instance, Davidson says: "Nagel is understandably astonished that a priori reasoning could show that our general picture of the world around us 'covering vast tracts of history, natural science, and ordinary lore' is largely true" (1999, p. 209). In "The Method of Truth in Metaphysics", he speaks of "a picture of the world that must, in its large features, be true" (2001a, p. 199). In both instances, it is clear that he takes it that a subject's global representation itself admits of veridical appraisal (see Chapter Three: Davidson's Holism for further discussion).

The concept of a global representation of the world has also featured prominently in recent discussions of coherence as a criterion or theory of justification. Bonjour (advocating a position he has since repudiated) argues that:

[a] system of beliefs might gradually *converge* on some definite view of the world and thereafter remain relatively stable, reflecting only those changes (such as the passage of time and the changes associated with it) which are allowed or even required by the general picture of the world thus presented....[and it is only in this] sort of case – the case in which the belief system converges on and eventually presents a relatively stable long-run picture of the world, thus achieving coherence over time as well as at particular times – that the coherence of the system provides any strong reason for thinking that the component beliefs are likely to be true. (1985, p. 170, italics in original)

Somewhat similarly, although with much less thoroughly coherentist commitments, Ernest Sosa uses the concept to explain how Descartes avoids vicious circularity in the *Meditations* (and how, according to Sosa, Moore, Quine and Davidson perhaps can). According to Sosa, "it is by adding interestingly to the coherence of one's picture of the world and one's place in it that one is able to gain a further measure of distinctive, epistemically valuable justification for one's empirical beliefs, a measure of justification that goes beyond the mere reliability of

those beliefs derivative from how we must acquire contents and form beliefs.” (2003, p. 182).⁸⁹

The concept of a global representation of the world also features prominently in the literature on explanation. Although it can be found within discussions of intentional explanation,¹⁰ it plays a more significant role within discussions of scientific explanation. Paul Humphreys says that it is “no explanation to provide a distorted representation of the world, and the ‘understanding’ induced by such models is illusory at best” (1989, p. 103). William Seager argues that “benignly emergent properties” although they are metaphysically inefficacious, feature in “useful, utterly indispensable, descriptive schemes, by which we bring order and understanding into our picture of the world” (unpublished, pp. 19-20). Jay Rosenberg also connects the idea of a global representation of the world to explanation and understanding, saying that understanding is (at least in many cases) “not a matter of analysis but of synthesis or, what is the same thing, of *explanation*, for to explain a phenomenon is precisely to bring it into systematic connection with other phenomena – to ‘subsume it under laws’, as the tradition has it – that is, to give it a ‘reading’ in terms of which it ‘fits’ into a larger, coherent, unitary picture of the world” (1981, p. 33). Todd Jones employs the concept in a paper addressed to the question of, as he puts it, “whether we can *epistemically* improve our picture of the world in ways other than gaining increasingly detailed reductive knowledge” (2003, p. 126, italics in original). And Gilbert Harman, remarking on a particular

⁸ Sosa elsewhere speaks about “a broader concern for the *kind* of knowledge we should seek in a good life. Wisdom might be one such, something closely connected with how to live well, individually and collectively. Another such might be a world-view that provides deep and broad understanding of major departments of human curiosity, which of course cries out for an account of what makes a curiosity proper.” (unpublished, p. 17 in pre-print, italics in original). The point Sosa here makes – that we would need an account of what makes a curiosity proper – is a fine one. What Sosa could have added, however, is that in addition to an account of this, we need an account of what a ‘world-view’ is.

⁹ See also Thomas Bartelborth, who connects the idea of a “plausible, detailed picture of the world” to that of having a “coherent belief-system in the sense of Bonjour” (1999, p. 211), and Stewart Cohen, who uses the concept to explicate a philosophical position (2002, p. 322).

¹⁰ See, for instance, Haack 2003, ch. 6, p. 17.

form of inference that “provides difficulties for most accounts of inductive inference”, says that “We can account for it if we assume, as I think we must, that inference to the best explanation is inference to the best total explanatory picture of the world” (1968, p. 533).

Again, these quotations are an illustrative sample rather than an exhaustive catalogue. In each case, the concept is put to work. Despite their philosophical differences, by employing the concept, Humphreys, Seager, Rosenberg, Jones and Harman seem to accept its coherence and usefulness and to share a general understanding of its character.

The literature on the origin, nature, fixation and change of belief also employs the concept of a global representation of the world. Perhaps most famously, Quine has claimed that “the stimulation of his sensory receptors is all the evidence anybody has had to go on, ultimately, in arriving at his picture of the world” (1969, pp. 75-6).¹¹ Michael Dummett, in the quotation that begins this chapter, takes it that the concept of a global representation of the world is essential to a proper understanding of judgment, saying that “A philosophical account of judgments lies, however, in their having a further significance than merely being right or wrong: they build up our picture of the world.” (1993, p. 157) The concept is also found within the literature on full versus probabilistic belief, as when Sven Ove Hansson uses it to explain the presence of full belief:

Our cognitive limitations are so severe that massive reductions from high probability to full belief (certainty) are inevitable in order to make us capable of reaching conclusions and making decisions. In other words, we treat things as certain although they are not. This rejection to full belief, or ‘fixation of belief’, helps us to achieve a cognitively manageable representation of the world.” (2003, pp. 39-40)¹²

¹¹ It is worth noting that Quine appeals to something like the unity of belief in numerous places. Consider a typical remark, from *Web of Belief*: “When an observation shows that a system of beliefs must be overhauled, it leaves us to choose which of those interlocking beliefs to revise; this important fact has come up repeatedly. The beliefs face the tribunal of observation not singly but in a body.” (1970, p. 13) We have here four notions – that of a web of belief, of a system of belief, of interlocking beliefs, and of a body of belief – where there is, in some sense, one thing that consists of or is made up of many things. Quine does not make clear what this unity consists in. Again, Chapter Three: Davidson’s Holism is relevant.

¹² See also van Fraassen 2001, p. 166: “it seems to be that there is a good place within the epistemic enterprise for having *one* picture of which you just say ‘that is the way things are’. Within that picture there remain many

It is also found within the literature on belief change, as when Anthony Gillies writes:

The best known, and most influential, theory of belief change is the AGM¹³ theory of belief revision.... The core intuition behind the AGM theory is that when forced to correct our imperfect picture of the world, we should only make changes to our beliefs which we are absolutely required to make. (2004, p. 3)¹⁴

But the concept of a global representation of the world is most prominent, within the literature on belief, in discussions which focus on the aim of belief. Sometimes, this is in the debate over whether true belief is more valuable than knowledge.¹⁵ More often, however, the concept is invoked, without much explication, in an account of the goal of inquiry. Stephen Luper argues that “a complete and accurate understanding of what is the case” is the goal of cognition, where this is a “view” or “picture” of what is the case (1990, p. 43). He, like Bonjour and Davidson, explicitly endorses the view that the picture can be accurate and that this accuracy comes in degrees. Somewhat similarly, Wayne Riggs argues that “the epistemic good life consists of having a comprehensive representation of the world that renders it intelligible to us” (Unpublished, p. 2). Riggs too takes it that this representation can be appraised in terms of its “accuracy”. van Fraassen argues that “The point of having beliefs is to construct a *single* (though in general incomplete) picture of what things are like...to form a single, unequivocally endorsed picture of what things are like” (1995, pp. 349-50).¹⁶ Jonathan

possibilities.” (Italics his.)

¹³ The AGM theory of belief revision draws its name from the first initials of the surnames of Carlos Alchourrón, Peter Gärdenfors, and David Makinson, the authors of “On the logic of theory change: partial meet contraction and revision functions”, *The Journal of Symbolic Logic* 50 (1985) 510-530, which set forth the theory.

¹⁴ See also Gillies 2003 (1,3,12,18), Gillies 2006 (118), and Chopra et al. 2003 (p. 422).

¹⁵ Patrick Hawley, for instance, employs the concept to argue that true belief is no more valuable than knowledge. “When deciding what to do, it is useful to have an accurate picture of the world. And when acting, it is useful to *maintain* an accurate picture of the world. But my picture of the world, with respect to proposition *p*, is no more accurate when I know *p*, than when I truly believe that *p*. And when maintaining an accurate picture of the world, knowing that *p*, and having a lasting true belief that *p*, are on a par.” (2003, p. 5)

¹⁶ van Fraassen also invokes the concept to explain the putative goal of science: “The ideal of unified science implies that the point of engaging in science is to construct a single (though at every humanly reachable stage incomplete) picture of what things are like. The suggests that the product of science is a set of propositions. Their intersection is the proposition which captures exactly that single picture of the world which science provides.” (1994, p. 339)

Roorda concurs, saying that “our beliefs are not just isolated sentences in a collection; they are meant to hang together, to tell a univocal story about the way the world is” (1997, 148). Richard Wollheim shares this view, saying that “the role of belief is to provide the creature with a picture of the world it inhabits” (1999, p. 13). Again, Wollheim takes this representation to be subject to global appraisal: “Not, of course, any picture of the world, but, subject to one proviso, a picture that depicts the world more or less as it is.” (1999, p. 13)¹⁷ In a discussion of Peirce’s view that ‘settled belief’ is the goal of inquiry, Cheryl Misak says that Pierce recognizes “that it will be suggested that we aim for a different kind of truth, something like the true complete picture of the world” (1987, p. 262). Simon Evnine, in a paper examining the believing of conjunctions, says that “Although humans, as limited and imperfect creatures, may well find within themselves pockets of belief isolated from the rational effects of the rest of their beliefs, this cannot be a healthy or acceptable situation. The goal is surely to arrive at a single, rationally integrated representation of the world (or as much of one as we can).” (1999, p. 221) And Michael Bratman says it is “part and parcel of our ordinary conception of belief” that one “should be able to agglomerate one’s various beliefs into a larger, overall view; and this larger view should satisfy demands for consistency and coherency” (1992, p. 4).

The concept of a global representation is also found within the philosophy of mind, although again this is typically as something used to explain something else, rather than as something that is itself the subject of examination. Vogeley et al. see the concept as central to a proper understanding of consciousness (1999, pp. 344-5), Papineau uses it to dispute a view of phenomenal concepts (2003), and Shoemaker explicitly connects it to the unity of

¹⁷ The proviso Wollheim has in mind is this: “there is also pressure within the creature to have some picture of the world rather than none. In many circumstances of life, it needs must prefer error to suspended judgement.”

consciousness (1996, pp. 184, 186). Perhaps most significantly, Andrew Brook argues that the concept is central to Kant's philosophy of mind,¹⁸ and that although cognitive science has taken up several Kantian insights (most notably concerning the unity of consciousness and the binding problem), it has neglected this aspect of Kant's philosophy of mind (1994).

The concept is also found, perhaps surprisingly, within computer science and artificial intelligence research. It is common to hear, of a robot, that "it instantaneously revises its beliefs to update its representation of the world" (Alechina and Logan, 2001, p. 102). Katherine Nelson and Richard Nelson invoke the metaphor to explain the difference between traditional and connectionist artificial intelligence (2002, p. 722). More significantly for philosophical purposes, it often arises in discussions of the frame problem, one of the key difficulties artificial intelligence researchers face. J Waskan and William Bechtel say that "since its introduction, the frame problem has been generalized as the problem of how a system could maintain an accurate representation of the world in which it acts" (1997, p. 34). John Pollock's more detailed description is:

The frame problem is a well-known but often misunderstood problem in the foundations of artificial intelligence. In the real world, things change. A rational agent residing in a changing world must be able to reason about change and persistence. This requires four kinds of reasoning. First, it must be able to dip into the world perceptually, acquiring information about the current state of the world. Second, it must be able to combine information obtained from different perceptual excursions and, through inference, construct a coherent picture of the world. Third, it must be able to detect changes perceptually and update its picture of the world accordingly. Fourth, it must be able to acquire general causal information about 'how the world works' and use that to predict the result of changes either observed by the agent or wrought by the agent's own actions. Officially, the Frame Problem arises only in connection with the fourth constituent of reasoning about persistence and change, but the solution to the Frame Problem requires principles that bear upon the other constituents as well. (1997a, p. 143)

The problem is due to McCarthy and Hayes 1969, a paper in which they claim that "A computer program capable of acting intelligently in the world must have a general

¹⁸ See also Baldner 1996, especially pp. 337, 341-2.

representation of the world in terms of which its inputs are interpreted” (p. 463). And Dennett describes the frame problem “as arising most naturally and inevitably as the problem of finding a useful, compact representation of the world” (1998, p. 207).¹⁹

Global Representation of the World

There are three interrelated concepts I hope to better understand: global representation of the world, overall global representation, and global representation of x , where x is a possible object of thought. The most perspicacious way of introducing these concepts will be to focus first on outlining in a bit more detail the concept of a global representation of the world, and then, once this concept is reasonably set out, use it to introduce the concept of overall global representation and global representation of x . I take this approach because I think the concept of a global representation of the world is already part of our conceptual repertoire. It should, therefore, be relatively easy to convey the broad outlines of the concept, and afterward use it to introduce the other notions. I should make clear that my characterization of all three concepts will be largely intuitive. I take it that this is the best way, or perhaps the only way, to proceed, as at this stage what I am trying to do is isolate concepts that I wish to better understand (through this dissertation, and research beyond). In other words, the characterizations of these concepts are largely intuitive because the concepts themselves are largely pre-theoretical, and it would be a mistake to build into the

¹⁹ Within the philosophical literature, see also Churchland and Churchland 1983, p. 14; Churchland and Sejnowski 1999, pp. 343-344; Dennett 1987, 32; Flanagan 1988, p. 546; Gibbons 2001, p. 27; Joyce 2004; Kornblith 1993, p. 105; Mandik and Clark 2002; Putnam 1998, p. 107; Ramsey 1990, p. 249; Sainsbury 1986, p. 104; Seager 2001; van Fraassen 1985, p. 255. For further uses of the concept in artificial intelligence research and other areas outside of philosophy proper, see also: Axel 2004; Bickerton 1990; Brazier et al., 1987; Brooks 1991; Dragani and Nissan 2004, p. 286; Garcia et al., 2000, p. 6; Guinn 1998; Harris 2001, pp. 247-8; MacKay 1969, p. 112; Mizzarro 2001, p. 2; Morganti 2004, p. 88; Pollock 1999, p. 79; Pollock 1997b, p. 526; Ram 1993; Rosenberg and Anderson unpublished; Smith, Hermelin and Tsimpli 2003; Sperber and Wilson 1995, p. 48; Zeevat 2003, p. 180.

intuitive notion the kinds of specific details that only an account of the phenomena falling under the concept could reveal.

In the balance of this section, I outline the concept of a global representation of the world, and focus in particular on giving a sense of what is philosophically puzzling about the claim that a subject's beliefs compose a *single* representation. Following this, I use the concept of global representation of the world to introduce the concepts of overall global representation and global representation of x . Finally, at the close of the chapter, I outline a number of distinct issues and problems and sharpen somewhat the questions I hope to make headway on in the remainder of the dissertation.

The Intuitive Concept

The first thing to say about the concept 'global representation of the world' is that this idea, or something very close to it, is often expressed via the locution "picture of the world". This expression is used in two senses, only one of which matches the sense of 'global representation of the world'. At times one hears philosophers speak of "the physicalist picture of the world" or "the realist picture of the world", where the idea is that the expression "picture of the world" picks out only certain broad features – in the 'physicalist picture', that everything is physical, and in the 'realist picture', that the world is mind-independent. In this use, people can be said to share a picture of the world even if their belief systems differ sharply. Two philosophers might share a physicalist picture of the world, for instance, even if they disagree about a great many things – about which physicalist theory of the mind is true, about whether Reagan was a better president than Clinton, about how humans are related to dinosaurs, or where the Gardasee is in relation to Switzerland, and so on. A "picture of the world" in this sense is one that captures only certain broad

features, or a fundamental general background within which disagreement about various details is possible, or even, perhaps, a methodological or philosophical stance of some kind.

This sense of picture of the world is not the one that interests me. The sense in which the expression is used that closely approximates the notion I am after is one on which a subject has a picture of the world that is, as it were, built up from all of her beliefs. On this notion, if S believes $p_1 \dots p_n$, then S's picture of the world is one according to which $p_1 \dots p_n$ are true, and S's picture of the world thus differs from the picture of the world had by any cognizer who either (i) fails to believe one or more of $p_1 \dots p_n$, or (ii) believes a proposition that falls outside this set.²⁰ The difference between A's picture of the world and B's picture of the world can be great or small, depending on how great the differences are in their belief sets. Your picture of the world is presumably much more similar to mine than it is to a frog's; you and I see the world in much the same way, whereas (presumably) a frog sees it much differently.

This sense of "picture of the world" is reasonably common (as demonstrated in the previous section), and closely approximates the notion I am after. That said, the expression "picture of the world" is deceptively obscure and potentially misleading. It is deceptively obscure in that it is so familiar that we might fail to notice that it is quite unclear what the expression even means. What is a picture of the world? What makes it a picture (why, that is, is the picture metaphor appropriate?), and what makes it of the world (rather than of

²⁰ I here ignore certain important issues, including that of indexical propositions, cognitive significance, and theories of propositions. I should also note that there might be weaker notions of 'same' in this context that are philosophically interesting. Two subjects would share a picture of the world, in the strictest sense, only if they had the same beliefs. But they might share a picture of the world in a weaker sense by virtue of having belief systems that overlap to a significant degree. An interesting question that I won't explore here is whether there is an important sense of 'same picture of the world' such that, necessarily, people who speak the same language share the same picture of the world. See however Chapter Three: Davidson's Holism for discussion of related issues.

whatever is wholly represented within the picture)?²¹ And it is potentially misleading in that the visual metaphor might tempt one toward unproductive avenues of understanding. The central problems are that a picture is, at least ordinarily, visual, and perspectival.²² If one has a picture of the world in the sense that interests me, that picture is clearly not visual. The danger is that the expression ‘picture of the world’ suggests these visual connotations. Of course this would never lead anyone to think that one’s picture of the world is in fact visual, but the power of the visual metaphor could close us to fruitful ways of understanding the relevant issues. Furthermore, pictures are perspectival in a way that some representations – such as models – perhaps are not. It may well be that global representations of the world are inherently perspectival, but if they are this is a view one should reach through investigation, rather than one we assume through the terms with which we frame the investigation. What I shall do, then, is minimize my use of the expression ‘picture of the world’, despite its familiarity and intuitive appeal, favoring instead the locution “global representation of the world”. The hope is that this expression is neutral, or at least more neutral, in terms of what it suggests about the structure and character of the representation.

There is, however, one way in which the metaphor of a picture of the world captures something essential about the intuitive notion of a global representation of the world. That is that such a representation is a unity of a multitude. Our minds, it seems, contain countless beliefs, these beliefs are about entities as varied as concrete objects, persons, places, properties, events, abstract objects, kinds, states of affairs, relations, and counterfactual possibilities, and the beliefs themselves can be either conscious or unconscious. Yet our

²¹ This last question has a parallel with actual pictures. Consider anything you would call a photograph of Providence. What makes it the case that this is a picture of Providence, rather than of merely whichever houses, streets, etc. fall completely within the photograph? More generally, the problem is this: if x is a picture or representation of y , what makes it the case that x is a picture or representation of y rather than of w or of z , where w is a proper part of y and y is a proper part of z ?

²² See, however, Lopes 1997 for an interesting argument to the effect that pictures are not necessarily visual.

minds may also be such that they have or contain an overall representation of reality that is, in some not well understood sense, constructed out of these individual beliefs. The picture metaphor vividly captures this. Consider a painting of the signing of the American Declaration of Independence, which contains or consists of numerous distinct representations – there is an area of paint that represents Thomas Jefferson, others that represent the various other signatories, their clothing, the tables and other furniture, the room in which the event took place, and so on. In addition to each of these individual representations, there is also the representation they together comprise. It is not merely a collection of representations, one of Thomas Jefferson, one of John Adams, one of a chair, one of a curtain, and so on; it is a representation of a complex state of affairs of which these are all features.

A central problem with understanding the concept of a global representation – perhaps the central problem – is understanding what the unity of belief consists in. If a subject indeed has a global representation, or picture, of the world, what makes it the case that she has a global representation rather than merely a multitude of individual beliefs? By virtue of what are her many distinct beliefs *a unity*?

We can get some sense of the problem by comparing a painting of the signing of the American Declaration of Independence with a collection of paintings, each of which represents one of the constituents of the event that was the signing of the American Declaration of Independence. The single painting does, and the collection of paintings does not, represent the signing. There are two reasons a collection of paintings – one of Thomas Jefferson, one of the chair, one of the room, etc. – is not a representation of the signing of the Declaration of Independence. The first is that the collection captures the constituents of the event, but not the relations in which they stand (relations which are necessary for the

event's being the event it is rather than some other). The second reason is that this collection is not a representation; each painting represents something, but the collection does not jointly represent anything, for the collection is not itself a representation. Suppose it were possible to represent in a painting (or set of paintings) that the constituents stood in the relations they did without representing in that painting (or set of paintings) the constituents; in this case, adding this additional painting (or set of paintings) to one's collection of paintings of the event's constituents would not yield a representation of the signing of the Declaration of Independence, for this new collection of representations that together represent everything there is to represent about the signing of the Declaration of Independence would not be unified – there would be no representation that represents the event itself, just representations that represent aspects of it. To represent a whole, one has to represent all its parts, and all the relations in which these parts stand, but the representations of these parts and relations must also be unified – they must together be one representation. One of my chief interests, and one that I make only limited headway on understanding in this dissertation, is that of what grounds this unity in the case of belief.

Another element of the pre-theoretical concept 'global representation of the world' is that, in addition to being based on, or built-up from, one's individual beliefs, a cognizer's global representation of the world is such that it is subject to global appraisal, where that includes three key points (i) the appraisal attaches to the representation of the world, rather than to the individual beliefs that comprise it, (ii) the appraisal concerns the quality of the representation *qua* representation, and (iii) the appraisal comes in degrees. A little more needs to be said about each point.

The intuitive notion of a global representation of the world is that it can be better or worse *qua* representation, where this is understood as a question of fit or faithfulness

between the representation and the world. The fit or faithfulness in question – that which holds between a representation of the world comprised of all of one’s beliefs, and the world – is distinct from the sum or average of the fit or faithfulness which holds between one’s individual beliefs and those aspects of the world they concern. (This is not to deny that there is a relation between the two levels of appraisal – that much is obvious.) To get a better sense of what I mean, let us consider an example that is somewhat easier to think about. I have, it seems to me, an overall representation of Brown University. It is impossible to say everything there is to say about the content of this representation within a dissertation, but for starters, the representation I have of Brown University is such that, among other things:

Brown University is named ‘Brown University’.
 Brown is in Providence, Rhode Island.
 Brown has an undergraduate college, a graduate school, a medical school, and no law school.
 Brown is not brown, although some of its buildings are.
 Brown’s president is Ruth Simmons; the immediate past president was a man who wore a bowtie and left for Vanderbilt because (he said) he didn’t feel he fit at Brown.
 The grandfather clock in the Registrar’s Office was given to Nicholas Brown by a man who ran slave trading ships on which about 50 percent of the captive slaves died from maltreatment.
 Brown has a main green and at least three smaller quadrangles.
 Sometimes the grass near the base of the trees on the main green is dead, probably from lack of sunlight.
 Brown’s post office closes earlier in summer than during the academic term.
 Brown has a College Curriculum Council that more than once has decided against introducing pluses and minuses into the grading system.
 Brown is not identical to any number, be it natural, irrational, imaginary, or otherwise.
 Brown is not shaped like a chiliagon.
 And so on....perhaps ad infinitum.

I also had, when I visited the university as a prospective graduate student in 1999, an overall representation of Brown. Naturally I cannot recall in detail the content of that representation, but I have no doubt that the representation I have of Brown now is better, *qua representation of Brown*, than the representation I had of Brown then. Exactly what this

means is hard to say (the issue is more difficult than it may first appear to be).²³ What is relevant for present purposes is that it seems undeniable that the representation I have now is better *qua* representation of Brown than the representation I had then, that this is not because (or not merely because) I had false beliefs then whereas I have true beliefs now, and that the notion of better is one that comes in degrees – no doubt the representation I had of Brown in 2001 was better than that I had of Brown in 1999, but not nearly as good as that I have of Brown now. The judgment or appraisal attaches to my representation of Brown and not to any of the individual beliefs that comprise it, it concerns the quality of the representation *qua* representation, and it comes in degrees.

The pre-theoretical notion of global representation of the world is one on which a global representation of the world is a unity of one's beliefs, and one on which a global representation is subject to appraisal *qua* representation. The pre-theoretical notion is largely silent, however, on the two questions that immediately fall out from this. The first concerns the exact relation between a subject's individual beliefs and her global representation of the world. It might be too generous to say that the pre-theoretical notion includes the idea that a subject's global representation of the world supervenes on her individual beliefs, for supervenience is a theoretical notion *par excellence*. But the pre-theoretical notion does include the idea that a cognizer's global representation of the world is, in some sense, built up from her beliefs, where this notion, when analyzed, would almost certainly involve a supervenience claim. The second question concerns what, exactly, it is for one representation of the world to be better, *qua* representation of the world, than another. The intuitive notion makes only vague gestures toward accuracy, or comprehensiveness, or

²³One might think the answer is simply that the representation I have of Brown now is made up of more true beliefs than the representation I had of Brown then, but this view, while perhaps ultimately correct, faces significant challenges (which I discuss in "The Measure of Knowledge", in progress). In a nutshell, the challenges centre on whether it is intelligible to understand how much a subject believes by appeal to enumeration.

completeness, but each of these terms stands in need of clarification. This is not to deny that these are legitimate terms of appraisal; consider again the differences between your representation of the world and that of a frog. The point is only that criteria for the application of these terms of appraisal, as well as a more thorough understanding of them, are needed.

Overall Global Representation and Global Representation of x

Having outlined in a bit more detail the concept of a global representation of the world, let me use it to introduce two further concepts, that of an overall global representation and that of a global representation of x . A subject's overall global representation is a subject's single, overall representation, whatever picture her beliefs, as a totality, present her with. A subject's overall global representation of x is a subject's single, overall representation of x , whatever picture of x her beliefs present her with. Let me say a bit more about each of these ideas.

It may be difficult to see the difference between the concept of a global representation of the world and the concept of an overall global representation. Here is the difference: the claim that a subject has a global representation of the world is the claim that a subject has an overall global representation that takes the world as its representational target. That is, there are two distinct elements to the concept of a global representation of the world. The first is that such a representation is a unity of all of the subject's beliefs. The second is that this representation takes as its target, or is a representation of, the world. An analogy will help convey the difference. Suppose each belief is an individual puzzle piece. A subject's overall global representation is whatever picture results from putting all these pieces together. A subject's global representation of the world is this picture whenever it purports to be a picture of the world. It is my view that a subject's overall global representation is,

necessarily, a global representation of the world, but this could be contested and hence it is important to distinguish conceptually between them.

A subject's global representation of x is a subject's single, overall representation of x , where x is a possible object of thought. Earlier in this chapter I spoke of the representation I have of Brown University, in a discussion in which I claimed that I have a representation of Brown University and that this representation is better, qua representation, than the representation I had of Brown University in 1999. What I was speaking of there is my global representation of x , where $x = \text{Brown University}$.

There are some difficult issues here that make it impossible, at this stage, to define more precisely the notion of a global representation of x . Here, in a nutshell, is the problem. It seems to me, intuitively or introspectively, that I have something like an overall representation of all kinds of things – of Brown University, of my mother, of the mind-body problem, of the computer on which I'm typing this dissertation, of snow, of the Second World War, and so on. This leads me to think I can isolate something like the concept 'global representation of x ', where x would be filled, respectively, by Brown University, my mother, the mind-body problem, the computer on which I'm typing this dissertation, and so on. At this point, however, familiar Fregean worries arise. Consider Lois Lane. Intuitively, she has a global representation of Clark Kent and a global representation of Superman, in that all her beliefs about Clark Kent compose something like a picture of Clark Kent, and all of her beliefs about Superman compose something like a picture of Superman. But there is some force to thinking that Lois Lane does not have an overall representation of that entity that is both Clark Kent and Superman. Here the claim is not that she fails to have a perfect representation of that entity, in the sense of knowing everything there is to know about the entity. Rather, it is that the beliefs she has that we would normally describe as beliefs about

Clark Kent do not seem to fit together with the beliefs she has that we would normally describe as beliefs about Superman to compose a representation of that entity that is Superman/Clark Kent. For this reason, there is some difficulty in making sense of the notion of a global representation of x . Some of these issues I will address later, in Chapters Five and Six, although not in a way that is very satisfactory. At this point the thing to say, I think, is that there seems to be some sense in which I have a global representation of Brown University, of my mother, of the mind-body problem, and so on. The familiar Fregean problem makes it unclear how to specify the general concept here, but that is a problem for a theory of such global representations to address.

Issues and Problems

To conclude the chapter, I will briefly distinguish a number of questions I take to be of interest and indicate what I will focus on in the dissertation. First, restricting the discussion to the concept of a global representation of the world, we can distinguish descriptive and normative lines of inquiry. A normative line of inquiry would investigate whether, or under what conditions, a subject ought to have a global representation of the world. Some of the quotations above seemed to endorse a view according to which a global representation of the world is a kind of cognitive achievement, something toward which inquirers ought to aim. A descriptive line of inquiry focuses, instead, on questions concerning whether, or under what conditions, a subject in fact has a global representation of the world, and on what such a representation is. Included in this descriptive line of inquiry are questions about the exact relationship between a subject's many individual beliefs and her global representation of the world (should she have one). These descriptive questions belong largely to the philosophy of mind, and my focus will be almost exclusively on them.

At this point two overlapping problem areas present themselves as most worthy of exploration:

- (i) Given that we have a concept of a global representation of the world, but no clear account of what such a representation is, we can either eliminate the concept or refine our understanding of it. The first question, then, concerns whether we should eliminate this concept or make sense of it. I want to know, therefore, whether the concept can be made more intelligible, or transparent, and whether anything in fact falls under it. In a sentence: What is a global representation of the world, and is there any reason to believe any or all subjects have one?
- (ii) If a subject (if any possible subject) has a global representation of the world, we need to understand the relation between this representation and a subject's many individual beliefs. Common locutions according to which a subject's beliefs 'tie together', 'fit together', 'bind together', 'form', 'construct', etc. suggest a kind of ontological dependence between a subject's many beliefs and her overall representation, but are obviously inadequate. We just don't know what they mean. Hence: What, exactly, is the relationship between a subject's individual beliefs and her global representation of the world (should she have one)?²⁴

These questions, asked about the concept of a global representation of the world, connect closely with similar questions one could ask about an overall global representation and a global representation of x . We should want to know whether to retain, or eliminate, these concepts. And we should want to know, again, what relation holds between a subject's individual beliefs and her overall global representation, or her global representation of some object or domain.

I set forth these questions to indicate a hope rather than to make a promise. When I embarked on this project I aspired to develop a clear and compelling account of global representation. I am only now coming to understand what the problems are that such an account would need to address. As Wittgenstein aptly put it in his *Notebooks*:

You are looking into a fog and for that reason persuade yourself that the goal is already close. But the fog disperses and the goal is not yet in sight! (June 1, 1915)

²⁴ There are obvious parallels here with what Peter van Inwagen calls the Special Composition Question and the General Composition Question. See van Inwagen 1987 and 1990.

The issues concerning global representations are, I believe, connected in interesting ways to the vexing philosophical problems of subjectivity, semantic content, the nature of propositions, Frege's puzzle of identity and cognitive significance, deduction, reference, and truth. It is no easy task to find a way forward on the problem of global representation, for these other problems throw long shadows over every possible path. To make progress on these questions, I will start by looking for help: the next three chapters focus on exploring philosophical work that involves the notion of a global representation of the world, or that might reasonably be seen to bear directly upon the questions and problems I have just adumbrated.

SOME INITIAL OBJECTIONS

How clearly we understand this talk of total views is questionable in the extreme.

– Michael Williams, “The Myth of the System”

It is not entirely clear how we are to understand the separateness of beliefs.

– Michael Wedin, “Aristotle on the Firmness of the Principle of Non-Contradiction”

The previous chapter sketched the concept of a global representation of the world and pointed to the pervasive use of the concept within contemporary analytic philosophy. It also outlined some of the issues I hope to better understand, for the fact that the concept is familiar and is widely used with the assumption that something falls under it shows neither that the concept is intelligible nor that anyone in fact has a global representation of the world, at this or any other world. To help make progress on these issues, I turn in Chapters Three and Four to the work of Donald Davidson, Robert Stalnaker and David Lewis, for each appeals to something like the notion of an overall global representation. But first I want to devote a chapter to considering three skeptical challenges to the project of the dissertation. The main purpose of doing so is to clear away some initial objections one might have to the project. By considering these skeptical worries I hope also to be able to give a somewhat better sense of what I mean by global representation.

I begin by considering arguments advanced by Michael Williams against the claim that we have what he calls a ‘total view’, which I interpret to be substantially similar to what I

call a global representation of the world. Williams's arguments are the only arguments I can find that directly oppose the claim that a subject has a global representation of the world.¹ For this reason, it is important that I address them individually. However, as I will argue, Williams's arguments are actually opposed to certain strongly coherentist conceptions of a belief system, rather than to either the claim that a subject has a global representation of the world or to the intelligibility of the concept of a global representation itself. In other words, even if Williams's arguments are sound, their true target is certain model of a belief system, and not the claim I examine and defend in this work. The discussion of Williams's arguments is fruitful because it helps me distinguish the unity of belief from coherentist claims with which the unity of belief could easily be confused or conflated. Following the discussion of Williams, I examine whether the massive modularity thesis within cognitive science and evolutionary psychology gives reason to doubt that a cognizer has a global representation of the world; I argue it does not. I am motivated to consider this objection in part because it has often been put to me informally, and therefore it seems worth defusing. Again, however, I also think that seeing why the objection fails can help make more clear what is meant by a global representation of the world. Finally, I close the chapter by arguing that an examination of whether a cognizer has a global representation of the world can proceed independently of an inquiry into the cognitive architecture that would realize such a representation. The discussion here, I hope, helps make clear that the unity of belief is not a claim about the physical realizers of belief states, nor about the neural architecture that subserves belief. Again, the point is to make the notion of global representation more clear by distinguishing it from distinct ideas and claims.

¹ There are other philosophers who endorse theses that might entail that a subject does not have a global representation of the world, but Williams is unique in explicitly considering the issue of whether a subject has a 'total view'.

Michael Williams on Belief Systems

One philosopher who professes to doubt that we have a global representation of the world is the epistemologist Michael Williams. He does not put his thesis quite this way, but it seems a reasonable extrapolation of what he says in chapter eleven of *Problems of Knowledge*. That chapter, titled “The Myth of the System”, begins with the question: Do we have a total view? Williams says “How clearly we understand this talk of total views is questionable in the extreme” (2001, p. 128), and he is surely right about this. Williams then develops a series of arguments against the claim that our beliefs constitute a coherent, integrated system. Williams takes it that his arguments are decisive against the claim that a subject has what he calls a ‘total view’, and for this reason it is worth considering his challenge. Doing so will also help us to distinguish the concept of a global representation from the concept of a coherent belief system.

Before I consider Williams’s several arguments individually, I want to note two general points about his discussion of ‘total views’. First, Williams appears to equate the claim that a subject has a total view with the claim that a subject has a coherent, integrated belief system in something like the way described by coherentists in epistemology²; his criticisms often target the claim that a subject has a “belief system or total view”, as if these are the same thing. (It is not clear to me whether he thinks the terms ‘belief system’ and ‘total view’ express the same concept, or merely have the same extension, but in either case he clearly takes it that a subject has a total view if and only if the subject has a belief system.) It is far from clear, however, that this is right. It may be that if the notion of a belief system, in the strong sense envisioned by coherentists, is merely a figment of philosophical

²There are, perhaps needless to say, many competing accounts of what it is for a belief system to be coherent, or of what the ‘coheres with’ relation is. One influential proposal is Bonjour’s 1985, on which coherence is a matter of logical and probabilistic consistency, pervasive and strong inferential connections between beliefs, and the absence of unexplained anomalies and unconnected subsystems of belief (pp. 95,98).

imagination, then there cannot be ‘total views’, or global representations of the world. But this is far from obvious, and on the face of it there seems to be little reason to even suspect that it is true. This conflation is relevant because Williams’s arguments are directed against the claim that a subject has a belief system of a particular kind (i.e., of the kind envisioned by coherentists). Even if those arguments are successful, therefore, he owes us an argument that if there are no belief systems, there are no total views.

The second general point to make about Williams’s discussion of total views, which builds on what has just been said, concerns his focus on coherentism. Williams’s argument against the intelligibility or existence of belief systems arises as part of a more general project in epistemology, in which he is concerned to defeat coherentist accounts of epistemic justification. Coherentism, as a theory of epistemic justification, rests on the claim that there are belief systems, given that beliefs systems are what are said to be more or less coherent. Establishing, therefore, that there are no belief systems would suffice to establish that coherentism (at least as presently developed) is false. This is a fine approach for Williams to take, but his execution is somewhat imprecise, for although he advertises his arguments as effective against the claim that there are belief systems, they tend either (i) to be effective only against coherentist accounts of belief systems, or (ii) to show that there is some tension between coherentism as a theory of epistemic justification and the notion of a belief system. Neither of these does anything to show that a subject does not have a belief system, for the concept of a belief system is not unique to coherentists. Foundationalism, for instance, typically depends on the claim that a subject has a belief system – it is just that one’s beliefs form a system such that most of one’s beliefs rests upon and gain justification from resting upon some other beliefs. (Williams himself describes the foundationalist project in these terms in an earlier chapter of the same book, saying: “In this [foundationalist] picture, a

system of justified belief is like a building: there is a bottom level – a foundation – on which all the upper storeys stand” (2001, p. 82, italics mine). Moreover, Williams seems to think that even a foundationalist belief system provides a total view, for at the start of the book he describes the Cartesian project as one on which Descartes “aims to construct a view of the world and our knowledge of it from the ground up.” (2001, p. 3) The point here is that even if Williams's arguments are sound, they show only that either people don't have belief systems in the way that coherentists describe, or that coherentism as a theory of epistemic justification can't be squared with the coherentist account of belief systems. In a sense, therefore, Williams's arguments are not directly relevant to the issue at hand concerning global representation. However, I consider each argument individually because they are the only arguments I am aware of that are billed as being directed against the view that a subject has a “total view”.

Turning to Williams's specific arguments: I will sketch each argument, argue it is unsuccessful even as an argument against the existence of belief systems, and then argue that the argument also fails if we try to extend it to an argument against global representations.

The first argument Williams offers raises doubts about whether the notion of a belief system can properly account for the fact that there are different kinds of acceptance. He says:

There is something fishy about the coherentist image of a ‘web of belief’. On this picture, a proposition is either in the system or out of it. This seems unrealistic. We have different styles of acceptance. For example, we might accept a certain theory as a promising working hypothesis, though we would not say that we believe it to be true. Like much traditional epistemology, the coherence theory places too much emphasis on belief as an all-purpose attitude of acceptance. Belief is a rather special normative attitude: a relatively unrestricted form of commitment. Not everything we accept (in some circumstances, or for some purposes) is something we unqualifiedly believe. (2001, p. 128)

In saying that it is “unrealistic” to maintain that every proposition is “either in the system or out of it”, Williams is claiming something like this: the notion of a belief system, understood as the kind of thing into which only propositions fully believed enter, does not do a good job of capturing the way a subject accepts propositions, for there are many ways to accept a proposition other than by believing it. Two things can be said in response to this argument. First, there is no reason to think that a subject can have a belief system only if belief is the only kind of acceptance (or the only kind of acceptance the subject engages in). Suppose some subject does not have a single way of accepting propositions – some she believes, others she takes as working hypotheses without believing, etc. It does not follow from this that the subject does not have a system of beliefs, where belief is understood as the strongest form of acceptance. It may be that the coherentist story is implausible if it is true that subjects accept propositions in a number of different ways. But this has no bearing on whether the claim that a subject has a belief system is true or false. Second, nothing prevents an advocate of the claim that a subject has a belief system from saying that what she intends to claim is that a subject has an acceptance system. Some of the propositions that are the content of items within that system are believed, others are held as working hypotheses, and so on. This would not seem to do violence to our ordinary notion of belief system – or, at least, no more violence than that done by accepting that there are kinds of acceptance other than belief. Again, it may well be that the coherentist story collapses if the items within one’s belief system differ in terms of how they are accepted, but the truth of this would not entail that a subject fails to have a belief system. It would show only that coherentism, at least as typically developed, is false.

Let us consider this argument of Williams’s directed against the view that a subject has a global representation, rather than against the view that a subject has a belief system. Is

the claim that a subject has a global representation of the world at odds with the fact (let us grant) that subjects at least typically accept propositions in more than one way? Well, at most this would seem to show only that so long as a subject accepts propositions in more than one way, she fails to have a global representation of the world. This would leave intact the project of specifying more clearly what a global representation of the world is, and of understanding the relation between a subject's global representation (when she has one) and her individual beliefs. In other words, the project of the dissertation would remain.

Williams's point would then amount to a friendly one: it would help us understand the conditions a subject has to meet if she is to have a global representation of the world. That said, however, I do not see reason to think that the fact that a subject at least typically (or even always) accepts propositions in more than one way entails anything about whether she has a global representation of the world. Suppose some scientist accepts, but does not believe, that some ideal condition holds, and that she accepts this for pragmatic reasons of some kind. A philosopher sympathetic to the claim that she nonetheless has a global representation of the world could say one of two things. First, that the scientist does have a global representation of the world, it's just that some of the content of that representation is believed to be the case, whereas other parts of the content of that representation is merely accepted as being the case. Or second, that the scientist has a global representation, understood as a unified representation of all and only what she believes. More work would need to be done to know which of these two paths one sympathetic to global representations should take. The general point is that neither path seems hopeless.

A second argument Williams offers builds on the claim that there are different kinds of acceptance by construing kind of acceptance as relative to context. Williams says:

[I]f there really are different styles of acceptance, and if an item's acceptance status varies from context to context, there may be no yes-or-no answer to the question of whether a given item is in or out of the system." (2001, p. 128-9)

The idea here seems to be that it makes no sense to say, without qualification, that S believes p (or that S Φ s p , for various kinds of acceptance). Rather, one believes p (or Φ s p) only relative to some context. A scientist might believe p at t relative to context A and merely accept but not believe p at t relative to context B. If this is right, then, one might think, to say that S believes p at t is not to express a proposition with content. There is no such thing as believing at t , without qualification – there is only believing at t relative to context A, believing at t relative to context B, and so on. Hence it doesn't make sense to say that a subject has a belief system, understood as involving her being such that there is some set of propositions each of which she takes to be true.

This line of argument strikes me as fundamentally flawed. A first problem is that there is little force to thinking that subjects believe only relative to contexts. Consider your belief that Santa Claus is a fictional character, or that Brown University is in Rhode Island. Is there any force to thinking that you believe those things only relative to some context? Certainly we can imagine your situation changing such that you come to believe that Santa Claus is real, or doubt that Providence really is in Rhode Island. But what that shows is only the wholly unremarkable fact that what you believe changes as your situation changes, and not that you believe only relative to a context.

The claim that style of acceptance is always relative to context also does violence to our ordinary notions of belief and of acceptance. Here are two ways of making the point.

First, Williams accepts the distinction between belief and other kinds of acceptance. That is, he grants that acceptance is a broader condition than belief, and that one accepts some propositions that one does not believe. It seems to me, however, that the distinction

between belief and other kinds of acceptance just is the distinction between unqualified acceptance and acceptance relative to a context. Williams himself seems to recognize this, for he says:

Take the case of accepting something as a working hypothesis or for the sake of argument. In its appropriate context...it may be deeply entrenched without being, in any unqualified way, believed. Indeed, we may accept something as a working hypothesis even when our long-term purpose is to falsify it. (2001, p. 129)

Williams's point here is that a proposition can be deeply entrenched without being believed. He makes this point to show that style of acceptance does not correspond to degree of entrenchment. However, the point rests on the idea that non-believing acceptance is acceptance relative to a context, whereas belief is unqualified acceptance, acceptance independent of context. If style of acceptance varies by context, then the distinction between belief and non-believing acceptance collapses.

Second, the claim that belief is relative to context is unintelligible, for part of the concept of belief is that it is unqualified acceptance. Michael Bratman makes this point, saying that it is "part and parcel of our ordinary conception of belief" that subjects do not believe relative to contexts. As he puts it, "Reasonable belief is, in an important way, context independent: at any one time a reasonable agent normally either believes something (to degree n) or does not believe it (to that degree). She does not at the same time believe that p relative to one context but not relative to another" (1992, p. 3-4). Why should one agree with Bratman about this? Well, consider what it is to wonder whether p . One is clearly not wondering whether, relative only to some context, one should assent to or grant p ; that would be to wonder something, but not to wonder whether p . What one is wondering is whether p is really the case. Hence what one aspires to, in wondering whether p , is to have a

context-independent judgment, a 'yes' judgment iff p and a 'no' judgment iff $\sim p$. But what one is doing when one wonders whether p just is wondering whether to believe p .

For these reasons, it is very hard to see that Williams could be right when he says subjects believe only relative to contexts. But if he isn't right about this, then the argument that moves from this premise to the conclusion that a subject does not have a belief system fails. Likewise, any effort to extend his argument to show the unintelligibility of global representations would be unsuccessful.

A third line of argument pursued by Williams presses the line that our beliefs do not comprise *one* system, but several:

Coherence theorists often speak of a belief-system's being held together by 'inferential' or 'explanatory' connections. But the idea of such relations tying our beliefs into a single system is evidently fictional. I could lose an enormous amount of my scattered historical knowledge without much impact on what I know of physics or mathematics. Our 'belief system', if we should talk this way at all, shows a considerable degree of *modularity*: We don't have a single system, but an indefinite number of sub-systems that are at best loosely connected.....Developments in the physical sciences...may lead to serious revision in our historical views. But developments like this can be understood without invoking the idea of a total view. (2001, p. 130)

The conclusion Williams is pressing here is that a subject's beliefs comprise numerous subsystems but no overall system (and thus, as he takes it, no total view). This conclusion appears to be reached via the intermediate conclusion that some of a subject's beliefs (such as her beliefs about various historical matters and her beliefs about math) are not tied together by explanatory or inferential relations, where this conclusion appears to be supported by the fact that the beliefs in question can be had independent of one another.³

³ On an alternative reading of the argument Williams is presenting, he moves directly from the premise that he could lose various historical beliefs without that impacting his scientific or mathematical beliefs to the conclusion that this historical beliefs and his scientific and mathematical beliefs are not part of the same system. In response: Perhaps there are interpretations of the word 'system' in which A and B are parts of the same system only if, once co-instantiated, the continued instantiation of each depends on the continued instantiation of the other, but this is not an obvious reading of the concept. There are, for instance, many functional systems that have as components C1 and C2, even though either could be present as part of a similar functional system in the absence of the other. One component of my computer, for instance, is the new

The argument runs roughly like this:

. Premise: I could lose an enormous amount of my scattered historical beliefs without much impact on what I believe about physics or mathematics.⁴

Lemma: Therefore, my beliefs about scattered historical matters are not tied together by explanatory or inferential relations with my beliefs about physics or mathematics.

Concl.: Therefore, my beliefs comprise numerous subsystems but no overall system.

Let us grant Williams the first premise, and moreover a stronger version of it, on which he could lose an enormous amount of what he believes about math and physics without that having much impact on his historical beliefs. In other words, grant that there is no dependence in either direction between having the historical beliefs and having the mathematical beliefs. The problem with the argument is that the intermediate conclusion does not follow, at least without more argument. It seems that to get it to follow one would have to insert something like the following premise:

Beliefs A and B are not tied together by inferential or explanatory relations if a subject believing both A and B could retain belief A if she lost belief B and retain belief B if she lost belief A.

I do not see any reason to think this principle is true. It is a bit difficult to assess because it is deeply unclear what ‘tied together’ means. But even on a rudimentary, intuitive understanding of this notion, it is hard to see why being ‘tied together by inferential or explanatory relations’ is supposed to require metaphysically necessary co-instantiation.⁵

memory chip I just installed; the fact that I could take that chip out again “without much impact” on the other components does not show that the memory chip does not now comprise, along with the other components, a system. Belief systems, of course, may not be anything like computer systems, and so it may be that for anything to count as a belief system it would have to satisfy the very strong requirement on system-hood this principle expresses. But Williams owes us an argument to that effect, if that is his view.

⁴ Williams speaks of knowledge, not of belief, but at this point in the dialectic he is not concerned with knowledge per se but with belief (that is, his point is that one does not cease to believe what one believes about physics when one ceases to believe what one believes about scattered historical events, rather than that one ceases to know what one believes about physics when one ceases to know what one previously knew about scattered historical events). For this reason I have recast the argument in terms of belief rather than knowledge.

⁵ See, however, the discussion of holism of belief content in Chapter Three: Davidson’s Holism, in which I consider whether appeal to the Humean principle that there are no metaphysically necessary connections between distinct existences could, when combined with a radical holism of content, explain the unity of belief.

Perhaps what Williams is getting at is that some of a subject's beliefs appear to have nothing to do with other beliefs had by the subject. My belief, for instance, that I skipped breakfast this morning has nothing to do with my belief that electrons are very small. There do not appear to be any explanatory or inferential relations holding between these two beliefs. They are thus, one might think, not tied together by inferential or explanatory relations. This alternative reading of the argument, however, is too hasty. At best we get, I think, that the first belief (call it Belief A) is not *directly* tied together with the second belief (call it Belief B). But for all that, it could still be the case that Belief A and Belief B are tied together *indirectly*, and are part of one overall system. The idea here is that while A and B have nothing to do with each other, there may still be indirectly tied together by virtue of each being tied to beliefs that have something to do with each other. For instance, I believe that if I hadn't skipped breakfast, I would have consumed matter, and that if I had consumed matter, that would include some electrons. (Or whatever.) The point is that inferential, explanatory or conceptual relations between beliefs could tie one's beliefs into one system even if it isn't the case that every belief is directly tied by inferential, explanatory or conceptual relations to every other belief. Of course, it is still quite unclear how inferential or explanatory relations, or the 'to do with' relation gestured at above, tie beliefs together, or what it even means for beliefs to be tied together (or for a belief system to be 'held together'). But insofar as we're required to consider these issues by appeal to these dark notions, it doesn't seem that Williams's argument is successful.

Let us grant, however, the conclusion of the first step of Williams's argument – that our believing is modular – and focus on the second step, according to which this modularity entails that a subject's beliefs are not part of one overall system. I do not see how this step is supposed to work, for the fact that something is modular does not, as a general rule, entail

that it is not systematic. In fact, quite the opposite is true: It makes no sense to call *x* and *y* modules unless they are parts of some system. To resolve the question of whether modularity entails that a subject does not have a belief system would require more refinement of the notion of a belief system than we should undertake, given that the focus here is on global representation rather than belief systems. Let us also grant modularity, therefore, and ask whether this does anything to show that a subject does not have a global representation of the world? It would seem to show that her beliefs are not held together into a global representation of the world by explanatory or inferential relations, but that doesn't show that she does not have such a representation. By way of comparison, think again of the unity of consciousness. My experience of a toothache and my experience of a blue sky appear to be 'modular' if anything is – I can have each in the absence of the other, there are no explanatory, inferential or conceptual connections between them, they appear to be achieved via distinct perceptual mechanisms, and so on. But for all that they are still, it is very reasonable to think, tied together into one overall conscious state. I don't have merely two experiences; I have one experience of which the blue sky and the toothache are aspects. Similarly, even if we grant that there are not explanatory or inferential relations holding between all of a subject's beliefs, this does not show that those beliefs do not compose one representation.

A fourth line of argument Williams pursues arises from the difficulty we have individuating beliefs. He says:

No one has the faintest idea how *many* beliefs he has, or even how to go about counting them. This isn't just because we have so many beliefs that we wouldn't know where to begin, though this is perfectly true. Rather, we lack clear criteria for individuating beliefs – that is, saying when beliefs are the same and different – without which there is no possibility of counting. Asking how many beliefs I have is like asking how many drops of water there are in a bucket: who's to say? I believe that my dog is in the garden right now; do I *also* believe that he is not in the house,

not in the basement, not in Siberia? Or are these beliefs somehow included in the original belief? This is an odd question. With respect to beliefs, we do not normally have use for fine-grained criteria of individuation. Unless we become captured by the coherentist image of 'our beliefs' as a finely articulated, complexly interrelated network – a 'web of belief' – I doubt we shall ever have any use for them. (2001, p. 131)

The problem to which Williams points is real. It is by no means clear what a belief is. It is not merely that we do not understand what a belief is as something physical, although that is a problem too of course. It is that we do not even understand where one belief ends and another begins. It is hard to see, however, how this is supposed to generate an argument against the view that a subject has a belief system, or against the view that a subject has a global representation of the world. Perhaps the idea is that the notion of a belief system is the notion of a system of beliefs $B_1...B_n$, and therefore to make sense of the notion of a belief system we have to make sense beliefs as enumerable entities. The problem presumably is metaphysical rather than epistemic – it's not just that *we lack criteria* for individuating beliefs, it's that *there are no criteria* for individuating beliefs. If this is right, it would do heavy damage to the notion of a belief system insofar as that is to be understood as a system of beliefs $B_1.....B_n$. But it is not clear that the notion of a belief system, or of that of a global representation of the world, is committed to fine-grained principles of individuation for beliefs. (A Lewisian account of belief, for instance, might be an account on which a subject has a belief system but no unique, correct set of individual beliefs. See Chapter 4 for discussion.) Moreover, if one genuinely has doubts about the intelligibility of the idea of an individual belief, then one might be expected to be more rather than less sympathetic to the idea of global representations. After all, Williams grants that a subject believes: the complaint is only with the idea that what a subject believes can be broken down into finely individuated

individual beliefs. But if it can't be so broken down, then a subject's believing is in some more fundamental way holistic.

We should be very cautious, however, about rejecting the concept of an individual belief. Doing so would do heavy damage to any epistemological story that involved the concept 'belief', for if there are no criteria that individuate beliefs then there are no beliefs.⁶ Suppose, for instance, that someone proposes an analysis of knowledge according to which "S knows that *p*" is analyzed as:

S believes *p*, S is justified in believing *p*, and *p*.

An account like this requires, for its cogency, that there be some fact of the matter about whether S believes *p*. Of course, the epistemologist who proposes this theory of knowledge does not need to know, herself, anything about who believes what. However, she is proposing a theory of knowledge, one that, among other things, should be capable of application to the question of whether Williams knows that his dog is not in Siberia. If, however, there is no way of individuating Williams's beliefs such that it is possible to say (even in principle) whether Williams does or does not believe his dog is not in Siberia, then the proposed theory of knowledge cannot be right. Moreover, the individuation problem spills into other areas of philosophy, and it is hard to see how Williams's eliminativist solution could be avoided there. For instance, if what individuates beliefs is propositional content, then to say that beliefs can't be individuated is to say that propositions can't be individuated (or 'finely' individuated). But what does this even mean?

⁶ As Donald Davidson has pointed out, it would damage more than our epistemology, for belief "is central to all kinds of thought. If someone is glad that, or notices that, or remembers that, or knows that, the gun is loaded, then he must believe that the gun is loaded. Even to wonder whether the gun is loaded requires the belief, for example, that the gun is a weapon, that it is a more or less enduring object, and so on" (2001b, p. 156-7). Lynne Rudder Baker provides a good overview of what is at stake in chapter six of her *Saving Belief: A Critique of Physicalism* (1987).

In summary: Williams argues that a subject does not have a “belief system or total view”, but the arguments at best show that certain ways of thinking about what a belief system is cannot be right. They do nothing to show that a subject does not have a 'total view' in the sense of a global representation of the world, since the conception of a belief system Williams targets is simply not the same as the conception of a global representation of the world. We can pull together the discussion above to make somewhat more clear what is meant by a global representation of the world by showing what such a representation is not.

A coherent, integrated belief system of the kind targeted by Williams is a belief system the constitutive beliefs of which, jointly, meet what we can think of as broadly logical conditions. One influential account of the 'coheres with' relation is Bonjour's 1985 proposal, on which coherence is a matter of logical and probabilistic consistency, pervasive and strong inferential connections between beliefs, and the absence of unexplained anomalies and unconnected subsystems of belief (pp. 95, 98). We can see that the concept of a global representation of the world is distinct from the concept of a belief system in this sense by noting that it is an open question whether a global representation of the world must meet any or all of these conditions. The central idea behind a global representation – unlike that behind the idea of a coherent belief system – is that the subject, in having such a representation, has something like a single judgment or view concerning what is the case. What we want to understand is what it is to have such a single judgment or view. This may involve the subject's individual beliefs standing in some appropriate relation to one another, but if this is so, this would be what grounds the unity rather than what the unity is.

We can make this last remark more clear, perhaps, by reflecting on the claim, commonly asserted by coherentists but by no means uniquely by them, that a subject's beliefs are “tied together” by explanatory or inferential relations. It is unclear whether those

who assert this mean to claim merely that there are explanatory and inferential relations of various kinds between beliefs, or that these explanatory and inferential relations do something, in the sense of determine some metaphysical state; this is at least suggested by saying that such relations *tie together* beliefs. On the first reading, saying that beliefs are tied together by explanatory and inferential relations is to say merely that explanatory and inferential relations hold between beliefs. On the second reading, to say that beliefs are tied together by explanatory and inferential relations is to say that such relations do something, bring a kind of unity to the beliefs, in much the way one might think that a bundle of sticks tied together is one thing and not merely many things. If Williams is right that there are no explanatory or inferential relations between beliefs, this shows that a subject does not have a global representation by virtue of such relations. If Williams is mistaken and there are explanatory or inferential relations between all beliefs, however, we are no closer to the claim that subjects have global representations. In other words, even if we were to accept that our belief systems are as coherentists imagine, we would appear to be no closer to understanding the unity of belief. What would be missing is an argument that shows how such relations confer or determine unity. Merely using the term 'tie together' to refer to inferential or explanatory relations that hold between beliefs does nothing toward this.

This point is worth emphasizing. It is commonly said (by both coherentists and non-coherentists alike) that explanatory and inferential relations “tie together” beliefs. It is equally common to hear the term “unified” used to name or describe a belief system that is free from logical inconsistency. However, the fact that these terms – 'tied together' and 'unified' – connote oneness of some kind does nothing to show that beliefs between which explanatory and inferential relations hold, and beliefs free from logical inconsistency, should in any sense

be thought of as a unity – that is, as one thing. The complaint here is similar to the complaint made by Lewis against Armstrong's necessitarian account of laws of nature:

Whatever N [the relation alleged to hold between universals] may be, I cannot see how it could be absolutely impossible to have $N(F,G)$ and Fa without GaThe mystery is somewhat hidden by Armstrong's terminology. He uses 'necessitates' as a name for the lawmaking universal N; and who would be surprised to hear that if F 'necessitates' G and a has F, then a must have G? But I say that N deserves the name of 'necessitation' only if, somehow, it really can enter into the requisite necessary connections. It can't enter into them just by bearing a name, any more than one can have mighty biceps just by being called 'Armstrong'." (1983, p. 366).

In the same spirit, the fact that terms like 'tied together' and 'unified' are used to name or describe beliefs or belief systems that meet certain broadly logical or formal conditions should not lead us to think that meeting such conditions either is what it is to be unified, or is a necessary condition on being so unified.

In this work I by and large do not take up the question of whether having a coherent belief system in the sense envisioned by coherentists is either necessary or sufficient for having a global representation of the world. (There is a relatively brief discussion in Chapter Four of whether logical consistency is sufficient.) That would be a worthy, but different, project. Moreover, the concept of coherence is not well understood, which would make the task of examining in detail whether coherence is necessary or sufficient for unity as much an inquiry into what coherence is as into what global representations are. However valuable the accomplishment of such a task might be it is beyond the scope of the present work.

Global Representation and the Modularity of Mind

A second challenge that might arise at an initial stage insists that the massive modularity thesis within cognitive science and evolutionary psychology threatens the claim that a subject has a global representation of the world. I have raised a number of problems with Williams's

defense of the modularity of believing, but much work has been done in other fields to establish a massive modularity thesis. We should therefore consider whether the massive modularity thesis is decisive against the view that cognizers have a global representation of the world.⁷ In this section, I will outline the thesis in question and argue that it does not threaten the claim that a subject has a global representation of the world. The discussion will do more than defuse the objection; it will also help make clear that the unity of belief, in the sense I have in mind, is a claim about the nature of a subject's beliefs at a time, and not about the etiology of those beliefs or about the processing that leads to or sustains doxastic states.

One of the first defenders of the view that the mind is modular was Jerry Fodor, who argued in *The Modularity of Mind* (1983) that modular input (perception) and output (action) systems assist a central, nonmodular reasoning processor. On this view, it is worth noting, the mind is modular but not massively so. As Dan Sperber has noted:

Although this was probably not intended and has not been much noticed, 'modularity of mind' was a paradoxical title, for, according to Fodor, modularity is to be found only at the periphery of the mind....In its center and bulk, Fodor's mind is decidedly nonmodular. Conceptual processes – that is, thought proper – are presented as a holistic lump lacking points at which to carve. (1994, p. 39; quoted in Clarke 2004, p. 138)

A more thoroughly modular theory of mind is espoused by John Tooby and Leda Cosmides, on whose view:

Our cognitive architecture resembles a confederation of hundreds or thousands of functionally dedicated computers (often called modules) designed to solve adaptive problems endemic to our hunter-gatherer ancestors. Each of these devices has its own agenda and imposes its own exotic organization on different fragments of the world. There are specialized systems for grammar induction, for face recognition, for dead reckoning, for construing objects and for recognizing emotions from the face. There are mechanisms to detect animacy, eye direction, and cheating. There is a

⁷ I am also motivated to consider this issue because a number of philosophers have, in informal conversation, expressed the view that massive modularity threatens the claim that a subject has a global representation of the world.

“theory of mind” module...a variety of social inference modules...and a multitude of other elegant machines. (1995, p. xiv; quoted in Clarke 2004, p. 138)

At first glance, a “massive modularity” thesis like this might be seen to threaten the view that a subject has a global representation of the world. However, there are two reasons this is not so. The first is that it is unlikely that the mind is *completely* modular, while the second is that modularity largely concerns cognitive processes rather than states of belief or representation. Let us consider each point.

In the quotation from Cosmides and Tooby above, cognitive architecture is conceived of as thoroughly modular, as consisting of nothing but distinct systems. Jerry Fodor, however, has argued that this model of the mind cannot account for central features of human reasoning, such as abductive inference and appeals to simplicity and conservatism. Fodor puts the point concerning abduction this way:

A lot of cognitive inference appears to be abductive. If so, then a lot of cognitive architecture can't be modular. Whereas modules are ipso facto informationally encapsulated, it's true, practically by definition, that abductive inferences are sensitive to global properties of belief systems. (2000, p. 98)⁸

Much the same point holds for appeals to simplicity and conservatism in theory choice. When deciding what theory to adopt (or, more prosaically, when deciding whether to believe something), we often appeal, or so it seems, to considerations concerning how our total epistemic commitments will be affected by the changes. Will the theory (or belief) under consideration simplify or complicate our total epistemic commitments? And will it require huge or minor revisions in what we believe?⁹ Yet these questions concern global properties of our belief systems. Simplicity, as used here, is not a property of individual judgments, but of one's entire body of belief. Similarly, the cost in terms of total doxastic impact of

⁸ To say that modules are informationally encapsulated is to say that information available for processing by one module is not available for processing by another. See Fodor 2000, p. 62.

⁹ These questions are independent. It could be that adopting some new belief would require huge revisions in what one believes, yet result in a simpler set of commitments once the shake-down is complete.

adopting (or abandoning) a belief depends on which belief system the belief will be (or is) a part of.

Fodor's criticisms of complete modularity are accepted even by some ardent supporters of modularity. Murray Clarke, for instance, has recently defended Cosmides and Tooby from Fodor by claiming that contra Fodor their view is only that the mind is "largely composed" of modules, not that it is completely so. As he sees it:

[T]here may well be some aspects of mental architecture that are nonmodular because there is a place where, to some degree, as Fodor says, "it all comes together". If that were not possible, abductive or global inductive inference would not exist. It would seem to follow that the mind cannot be completely modular. (2004, p. 7)

The point here is that however much the mind may be modular, it cannot be completely so. Some cognitive processes are domain general rather than domain specific, and involve appeal to properties of the totality of one's doxastic commitments.

But what if the mind *were* completely modular? Even then, the massive modularity thesis would have little bearing on the question of whether a subject has a global representation of the world. This is because the massive modularity thesis is a thesis about cognitive processes not about the output of such processes. It is the claim that cognitive processes are functionally insulated from one another, rather than the claim that the doxastic states such processes produce are isolated from one another. As Stephen Pinker's metaphor describes the view, it is that "the human mind...[is] not a general-purpose computer but a collection of instincts adapted for solving evolutionarily significant problems – the mind as a Swiss Army knife." (1994, p. 420). The view that our cognitive processes deliver one overall representation is compatible with the view that cognitive processes are domain specific, or at least it is reasonable to assume so in the absence of evidence to the contrary. A house, for instance, is no less one thing for having been brought into existence by domain-specific

workers (the foundation diggers, the framers, the roofers, the drywallers). It *may* be that believing cannot be like this, that beliefs that result from domain specific processes are isolated from beliefs that result from other domain specific processes. But unless we have compelling reason to think this is so even the strong claim that *all* cognition is modular does not weigh against the claim that a subject has a global representation of the world, understood as something like one judgment concerning what is the case.

At this point in the dissertation I have not said much about what I take the unity of belief to consist in. The lesson of the preceding discussion is that in trying to understand the unity of belief, we should be clear that what is at issue is not a thesis about how the mind performs its cognitive functions. The question of whether there is unity or mere diversity there is a different question from that of whether a subject's beliefs are unified into one representation.

Global Representation and Cognitive Architecture

A third initial objection rests on the fact that if a subject has a global representation of the world, that representation must be realized in the subject's cognitive architecture. For this reason, one might think that an examination of the claim that a subject has a global representation of the world, which I undertake in this dissertation, cannot proceed independently of an investigation of cognitive architecture. After all, if, say, the correct theory of cognitive architecture is connectionism, and connectionism cannot account for a subject's having a global representation of the world, then a subject cannot have a global representation of the world. In other words, the fact that a global representation of the world would have to be realized in cognitive architecture might lead one to think that in order to understand what such a representation is, and establish whether a subject has such a

representation, one would have to investigate cognitive architecture and find there an appropriate realization base.

I think this challenge fails, and that it is possible to investigate the concept of a global representation of the world and establish whether a subject has such a representation without taking a stand on which theory of cognitive architecture is correct. Again, considering this initial worry will help give more shape to the concept of global representation and make the project of the dissertation more clear.

Why does the challenge fail? For three reasons. First, there is no consensus on which theory of cognitive architecture is true and thus there is no theory of cognitive architecture with which an account of global representation must be compatible if it is to be taken seriously. Second, if subjects do in fact have a global representation of the world, then that is a datum that a theory of cognitive architecture must accommodate; for this reason the question of whether a subject has a global representation must be settled prior to our giving final assent to some particular theory of cognitive architecture. Third, one of the questions that interests me is whether it is metaphysically necessary that a subject has a global representation of the world. Yet theories of cognitive architecture are generally thought to concern contingent structures. Let us look at each point in more detail.

The first reason we have for thinking that the current project can proceed without taking a stand on a theory of cognitive architecture is that there is no generally accepted view with which a theoretical account of global representations of the world could be required to be in accord. As Dan Sperber has said, “our understanding of cognitive architecture is way too poor, and the best we can do is try and speculate intelligently” (2002, p. 47). Even Jerry Fodor, who is one of the most prominent defenders of classical computationalism, one of the two main options on the table, has recently expressed serious doubts about whether his

favored theory is true, or even a large part of the truth concerning the mind. Fodor says, for instance, that he is “inclined to think that, sooner or later, we will *all* have to give up on the Turing story as a general account of how the mind works”. (2000, p. 47, italics his) Fodor’s worries arise from the fact that the computational model cannot account for cognition that involves appeal to global properties of our doxastic commitments:

I suppose that the moral will eventually be conceded; namely, that the Computational Theory is probably true at most of only the mind’s modular parts. And that a cognitive science that provides some insight into the part of the mind that isn’t modular may well have to be different, root and branch, from the kind of syntactical account that Turing’s insight inspired. (2000, p. 99)

Unfortunately, Fodor’s view of the computational theory’s main rival, connectionism, is even dimmer; as he puts it, “the standard current alternative to Turing architectures, namely connectionist networks, is simply hopeless.” (2000, p. 47) The upshot, at least by his lights, is that there are cognitive properties that no current theory of cognitive architecture has the resources to explain:

The substantive problem is to understand, even to a first approximation, *what* sort of architecture cognitive science ought to switch to insofar as the goal is to accommodate abduction. As far as I know, however, nobody has the slightest idea. (2000, p. 47, italics in original).

The point here is not merely that there is little reason to require a theoretical account of a phenomenon (here, that of a subject’s having a global representation of the world) to be in accord with some particular theory if that theory is not widely accepted. If that were the only claim, then it might still make sense to require that the thesis be in accord with at least one of the theories that are matters of controversy. The point is a stronger one: not only do we not know which theory of cognitive architecture is true, we do not have compelling reason to think that any of the known options must be. Indeed, if Fodor is right, we have reason to think that none is. For this reason, it is inappropriate to require that an investigation of

global representation be in accord with any of them, let alone with some particular one among them.¹⁰

The second reason it is unreasonable to require that, in investigating global representation, I develop or commit to some particular account of cognitive architecture, is that if a subject does in fact have a global representation of the world, this is a datum that a theory of cognitive architecture would have to accommodate. If it turns out that, for instance, some particular theory of cognitive architecture cannot account for our having a global representation of the world, and we do in fact have such a representation, then that theory is obviously false. An argument against our having a global representation of the world that includes as one of its premises the truth of that theory of cognitive architecture would be unsound. Similarly, if some theory of cognitive architecture accounts for or entails that each of us has a global representation of the world, and we do not in fact have any such representation, then an argument designed to establish that we have such a representation by asserting as a premise that such a theory of cognitive architecture is true obviously fails. For this reason, it is important to argue for or against the claim that subjects have global representations without employing arguments that rest on some particular theory of cognitive architecture.

The general point here is that theories of cognitive architecture are supposed to answer to the facts, as best we can make them out, concerning how the world is represented within thought. They are not supposed to adjudicate those facts. In order to know which theory of cognitive architecture is true, we must first have a good idea of what the data are that such theories are supposed to account for. The only way that I can see of developing

¹⁰ Of course, some proponents of one or another approach to cognitive architecture may think, unlike Fodor, that there is compelling reason to think that their approach is close to the full truth on the issue. But the fact that some people may think this does nothing to change the fact that this is not a consensus opinion.

this, however, is by relying on whatever arguments and insights we can find that are independent of the details of our current theories of cognitive architecture. We need to first determine which data a theory of cognitive architecture is supposed to accommodate before we can be in a position to know whether some particular theory is true; this is why we cannot give assent to some theory of cognitive architecture and then, using that theory as a premise, show that some cognitive phenomenon (such as our having a global representation of the world) fails to obtain.

Finally, one (epistemic) possibility I am interested in investigating is the claim that it is metaphysically necessary that a subject has a global representation of the world. Theories of cognitive architecture, in contrast, are generally thought to concern the contingent; they describe how a cognizer's representations happen to be realized. For this reason, I could not expect to find support for the modal claim within a theory of cognitive architecture, for the best such a theory could show is that we happen to have a global representation of the world. That is, suppose we had a theory of cognitive architecture that was well established. If that theory entailed that a subject's many beliefs are unified into one representation of the world, that might be thought to lend some support to the claim that we do, in fact, have a global representation of the world. But this would do nothing to show that cognizers must have a global representation of the world, which is a claim I wish to examine.¹¹

It is worth acknowledging that it might seem that connectionism is a natural ally of anyone wishing to argue that subjects have global representations, for "it is often plausible to view such [connectionist] networks as collectively or holistically encoding a set of propositions" (Ramsey, Stich, Garon, 1990, p. 511). On this picture, a subject does not have

¹¹ I agree as well with Braddon-Mitchell and Jackson when they say:

[U]nlike many realists in the philosophy of mind influenced by cognitive scientists, we do not think that the details of the inner workings – a fascinating and important topic in cognitive science – tell us about the *essence* of the mental. We think that beings with very different cognitive architectures might still have mental states of the same general kind as ours. (1996, p. x)

a number of discrete states each of which represents some proposition as being the case. However, the issue is not clear, for to say that a subject has a global representation of the world is to say only that the subject, somehow, has a single representation that includes as its content everything that the subject believes. It is not to say that the subject fails to have individual states that represent or encode some proper part of what she believes. Of course, it may be that a subject has something like a connectionist network, the content of which is a global representation of the world, and that what we call her individual beliefs are merely individual propositions that fall within the content of the one global representation. But it may just as well be that there are discrete representational states that are bound together into one representation. Here the issue is akin in some ways to the debate within the philosophy of mind on the unity of consciousness. One view, perhaps the standard view, is that a subject has at a time some large number of conscious representations or conscious experiences which are then bound into one unified field of consciousness. An alternative, championed by Michael Tye (2003), is that a subject has at a time just one conscious experience, where the putative multiplicity (blue here, a sound there, a fleeting thought here, etc.) is in fact representational content rather than a reflection of a multiplicity of representational vehicles (i.e., experiences). On this view, a subject has at a time exactly one experience, which is of n different 'things', rather than n experiences, each of which is of one thing. The global representation thesis is thus *prima facie* compatible with both the connectionist picture of holistic representation, and with what Stich (1983) has called propositional modularity, the view that propositional attitudes (including beliefs) "are *functionally discrete, semantically interpretable*, states that play a causal role in the production of other propositional attitudes, and ultimately in the production of behavior." (Ramsey, Stich

and Garon, 1990, 504, *italics theirs*) The issue is complex, to say the least, but if the arguments given above are right it is not one we need engage in at this point in the project.

This chapter has focused on considering three initial objections to the project of this dissertation. The aim has been both to defuse initial worries that might arise, and to use the objections constructively, to help bring the notion of global representation, and the claim that a subject's beliefs comprise a global representation of the world, into better focus by distinguishing them from distinct concepts and claims. Whatever ultimate connection there may be between the unity of belief and coherentism, however the latter concept should be understood, the claim that a subject's beliefs are unified is not the claim that a subject's beliefs meet certain conditions of consistency or coherency. (It may be that coherence is necessary for unity; that is an issue I set aside. The point is that in speaking of the unity of belief I mean by unity genuine unity – that is, oneness – and not something different such as consistency or coherency.) The oneness at issue, as consideration of the second objection suggests, is oneness in a subject's representation of the world. The subject does not merely represent this and that and something else: there is some one way she represents the world as being. The unity in question is not (or need not be) at the level of the cognitive mechanisms that generate beliefs. Finally, the issue I am interested in is not about the cognitive architecture that underlies our thinking and believing. It is rather about more general questions about how the world is represented within thought, by us, and perhaps by any subject.

I turn now to considering the work of Donald Davidson (in Chapter Three), and of Robert Stalnaker and David Lewis (in Chapter Four), for all three defend claims that seem to involve something like the notion of a global representation of the world. My aim will be to

see whether we can glean anything from their work that helps us understand the issues and problems outlined in Chapter One.

DAVIDSON'S HOLISM

The idea of defining a symbol in use was, as remarked, an advance over the impossible term-by-term empiricism of Locke and Hume. The statement, rather than the term, came with Frege to be recognized as the unit accountable to an empiricist critique. But what I am now urging is that even in taking the statement as unit we have drawn the grid too finely. The unit of empirical significance is the whole of science.

– Quine, “Two Dogmas of Empiricism”

Perhaps the philosopher who has done the most to press the idea that there is a certain kind of unity to belief is Donald Davidson. Much of his approach draws on earlier philosophers, notably Quine, Hempel and Wittgenstein, and aspects of his project intersect with issues pursued independently by Wilfrid Sellars, Gilbert Harman, Hartry Field, Ned Block, Robert Brandom, and others.¹ But Davidson’s treatment is broadest and among the most sustained and prominent, and for this reason it is worth considering his discussion of what he calls *the holism of the mental*.

A full examination of Davidson’s holism would take us too far afield – his treatment of the issue is bound up with a much broader project in the philosophy of language and mind, and involves a number of overlapping claims, not all of which are relevant to the goal of understanding the unity of belief. In what follows, I will focus on those aspects of Davidson’s project that have the clearest connection to the issues I hope to understand concerning global representations. My goal in this chapter is twofold: to investigate whether Davidson’s body of work involves, or concerns, the unity of belief and the concept of a

¹ I have in mind Quine 1951; Hempel 1950; Sellars 1963; Harman 1974; Field 1977; Block 1986; and Brandom 1994.

global representation, and to determine whether Davidson provides an account of the nature of such a representation, or of what it is in virtue of which a subject has such a representation (should she in fact have one). I should note that I am interested in determining not merely whether Davidson explicitly provides such an account, but whether his central insights concerning the interdependence of the mental can be used to generate such an account. I will argue, in regard to the first goal, that an important aspect of Davidson's work involves – and indeed can reasonably be thought to rely on – the concept of a global representation of the world. In regard to the second goal, I will argue that Davidson does not provide an account of the unity of belief and of global representation, and that his insights cannot be extended to provide such an account – or, at least, that my efforts to extend them to provide such an account fail.²

The chapter proceeds as follows. I discuss in turn three different theses defended by Davidson, which I label belief holism, holism of belief content, and interpretive holism. Neither of the first two theses involves the notion of a global representation, but I discuss them because they could reasonably be seen to be closely connected to the unity of belief and the issue of global representation. Davidson's discussion of the third thesis, interpretive holism, explicitly relies on the notion of a global representation of the world, and Davidson appears to be sensitive to some of the issues that I take to be most interesting. I will argue, however, that none of these three theses provides an adequate account of the unity of belief. The chapter will close with a brief discussion of Davidson's claim that beliefs are (or at least sometimes are) compartmentalized. This claim is, as he recognizes, in tension with the holism of the mental, but he finds it necessary to invoke the notion of compartmentalization

² I should note here that I do not take myself to be arguing against Davidson, for Davidson does not claim that he provides an account of the unity of belief or of global representations. Rather, my interest is in seeing whether such an account can be found within or developed from his work.

to explain irrationality. Consideration of this issue will help to illustrate that although Davidson appears to endorse something very close to the claim that a subject's beliefs compose a global representation of the world, he hedges or withdraws his assent from this claim in an important way.

What Thought Requires

A basic claim of Davidson's, which we can take as a starting point, is that to entertain any proposition one has to believe many other propositions. As he says:

If you wonder whether you are seeing a black snake, you must have an idea of what a snake is. You must believe things such as: a snake is an animal, it has no feet, it moves with sinuous movement, it is smaller than a mountain. If it is a black snake, then it is a snake and it is black. If it is black, it is not green. Since you wonder what you are seeing, you must know what seeing is: that it requires the use of the eyes, that you can see something without touching it, and so on. (2004a, p. 10)³

The idea here is that entertaining a proposition requires the deployment of concepts, and the possession of concepts involves having the ability to classify objects as falling under the concepts, which in turn involves having at least some true beliefs about the conditions something has to fulfill to fall under the concepts in question. As Davidson says elsewhere, “we would deny that someone had the concept of a man who did not know something about what distinguishes a man from a woman, who did not know that fathers are men, that every man has a father and a mother, and that normal adults have thoughts.” (2004c, pp. 136-7). Moreover, there are logical and semantic relations among concepts, such that to have certain concepts entails having certain other concepts. Of course, one might have

³ Davidson puts the point much the same way in “Thought and Talk”:

[B]elief is central to all kinds of thought. If someone is glad that, or notices that, or remembers that, or knows that, the gun is loaded, then he must believe the gun is loaded. Even to wonder whether the gun is loaded, or to speculate on the possibility that the gun is loaded, requires the belief, for example, that the gun is a weapon, that it is a more or less enduring physical object, and so on. There are good reasons for not insisting on any particular list of beliefs that are needed if a creature is to wonder whether a gun is loaded. Nevertheless, it is necessary that there be endless interlocked beliefs. The system of such beliefs identifies a thought by locating it in a logical and epistemic space.” (Davidson, 2001b, pp. 156-157)

doubts about the example Davidson provides, concerning what has to be the case if one is to wonder whether one is seeing a black snake. We can probably imagine possible situations in which it seems right to say that the subject in question wonders whether she is seeing a black snake, but who lacks all the beliefs Davidson mentions. And it is certainly possible to imagine situations in which she lacks some of the beliefs Davidson mentions. This is no objection to Davidson, however, for his point is only that a subject has to have many beliefs about snakes in order to wonder whether something is a snake, not that a subject has to have some canonical list of beliefs about snakes in order to wonder whether something is a snake.

As he says:

I do not wish to give the impression that there is a fixed list of things you must believe in order to wonder whether you are seeing a black snake. The *size* of the list is very large, if not infinite, but membership in the list is indefinite. What is clear is that without many of the sort of beliefs I have mentioned, you cannot entertain the proposition that you are seeing a black snake; you cannot believe or disbelieve the proposition, wish it were false, ask whether it is true, or demand that someone make it false. (2004a, pp. 10-11, italics in original)

The idea is that in any possible situation in which a subject wonders whether she is seeing a black snake, she must have some true beliefs about snakes sufficient to make it the case that she has, as Davidson says, “an idea of what a snake is”. If she lacks the beliefs Davidson mentions, that snakes are animals, that they have no feet, etc., perhaps she believes that snakes are long and thin, that they are often found under rocks, that some are venomous, that they shed their skins, and so on. Again, in Davidson’s view there is not some specific number of such beliefs that are required, nor some definite list. Rather, the idea is that there must be sufficient beliefs, where this could be satisfied in more than one way.

All we have considered so far is Davidson’s claim that to *entertain that p* involves having many (true) beliefs that *p*. From here, it is a short step to belief holism.

Belief Holism

‘Belief Holism’ is sometimes used to name the view I call ‘Holism of Belief Content’. In this chapter, I will use ‘Belief Holism’ in the following sense:

Belief Holism = necessarily, if a subject has any belief she has many other beliefs.

The argument for belief holism follows straightforwardly from the argument for what thought requires. Suppose Sally believes that she is seeing a black snake. This requires that she deploy the concepts “black snake” and “seeing” (perhaps among others, such as “now”, and some self-referring concept). But to have these concepts requires that one have many other beliefs about what black snakes are and about what seeing is. Having these beliefs, in turn, will involve some further concepts (various other things will be held true of blackness, of snakes, and of seeing). The possession of these concepts will depend on the having of further beliefs, involving novel concepts, and so on. To possess just a ‘simple’ belief, such as that one is seeing a black snake, requires on Davidson’s view an enormously rich conceptual system and, therefore, an enormously rich belief system.

Belief holism is widely but not universally accepted.⁴ I am not interested here in whether the claim is correct. Rather, the focus is on determining whether, even if correct, it provides, or can be extended to provide, an account of the nature of global representations. I will argue that although at first glance belief holism might seem promising in this regard, a negative verdict is called for.

⁴ Belief holism is usually thought to be a modern invention, but that there are entailment relations between the having of beliefs or concepts has been recognized since at least Plato. For instance, in *The Apology*, Socrates says:

Does any man, Meletus, believe in human activities who does not believe in humans?...Does any man who does not believe in horses believe in horsemen’s activities? Or in flute-playing activities but not in flute-players? No, my good sir, no man could. (Plato, *Apology*, 27b)

Socrates is pointing to entailment relations between the having of beliefs or concepts: necessarily, if S believes that there are horsemen’s activities, S believes that there are horses (or, alternatively: necessarily, if S has the concept ‘horsemen’s activities’, S has the concept ‘horses’).

Belief holism does not appear to involve any talk of a global representation; all that is claimed is that, necessarily a subject has many beliefs if she has any beliefs. However, one might still think belief holism is all that is needed to account for a subject's having a global representation of the world. The thought runs roughly like this: if a subject necessarily has many beliefs if she has any beliefs, then a subject necessarily has a whole system of belief any time she has any beliefs at all. Moreover, what would such a system be if not a subject's taking the world to be some particular way, which is tantamount to a subject's having a single, overall representation of the world? To try to develop this point, it might be helpful to put it into the mouth of an imaginary interlocutor:

Belief holism entails that a subject has a global representation of the world, and provides an account of why this is so. For it teaches us that beliefs are not, as Ernie Lepore has put it, "little atoms". A subject cannot have just a single belief, or even a handful of beliefs. She cannot merely believe that, for instance, she is seeing a black snake in the grass. Rather, she has to believe a large, rich picture according to which she is seeing a black snake in the grass, to see is to have eyes, to be a snake is to be an animal, being black precludes being completely green, and so on. She has to represent all of that as being the case if she is to represent any of it as being the case: hence the primary object of her attitudes is not a single, isolated proposition but a whole system of propositions.

On this position, although belief holism does not explicitly claim that a subject has a global representation of the world, it can help us make sense of a subject's having such a representation.

In my view, this is too quick, and fails to grasp the philosophical problem. It is relatively uncontroversial that all of us have many beliefs.⁵ That was never in question. What is hard to understand is why we should think – and what it is to think – that we have a single, overall representation that those beliefs somehow compose, and what relation this representation, should there be such a thing, stands in to one's (other) beliefs (what, that is,

⁵ But see Chapter Four: Lewis and Stalnaker, which considers David Lewis's claim that 'beliefs' is a bogus plural.

the compositional relation is). In other words, the philosophical problem *rests on* the fact that we all have many beliefs. If each of us had just one belief, then there would be no problem in understanding how each of us has a single overall representation. We would represent something as being the case – whatever the one thing is we believed – and that would be our single overall representation. Moreover, there would be no need to understand the apparently compositional relation between a person's single overall representation and her other beliefs, since she would have no other beliefs. (Both the singleness, and the overallness, would come cheap.) What belief holism adds to the truism that we all have many beliefs, the very truism that motivates or underpins the problem, is that it is necessary that we have many beliefs given that we have any. And not just us: anything that counts as a believer has many beliefs. But how does the fact that it is necessary that believers have many beliefs entail that, or explain why, a subject has a single, overall representation? Or, to put the question another way, how does its being necessary that believers have many beliefs make it the case that each believer has a single overall representation rather than merely a collection of many representations?

Belief holism, I will argue, does not give us reason to believe a subject's beliefs compose one representation, nor offer an account of why this is so. There are three versions of the argument, depending on the details of the belief holism under consideration. It is not clear which version of belief holism is closest to Davidson's, although he would likely accept all three versions in the sense of agreeing that the sentence that expresses the view states a true proposition.

Consider first the version of belief holism already articulated, one which holds only that a subject necessarily has many beliefs if she has any beliefs. (This version is distinguished from the two versions to follow because it says only that to believe any given

proposition one has to believe many other propositions, not that to believe any given proposition entails believing other, specific propositions.) Suppose at the actual world S believes, among many other things, that a black cat is on a mat. It follows from this version of belief holism that she could not lose all her other beliefs (without acquiring any new ones) while retaining the belief that a black cat is on a mat. However, this version of belief holism leaves open that she could lose any particular belief she has while retaining the belief that a black cat is on a mat. Consider, then, the possible world just like the actual world but in which the subject in question fails to have some belief that she has in the actual world, say the belief that it rained on Jan. 1, 1900 in New York. Call the actual world W and the world just like W except that the subject fails to believe that it rained in New York on Jan. 1, 1900 W^* . If a subject's beliefs form a single, overall representation, then S 's beliefs in W form a single overall representation and S 's beliefs in W^* form a single overall representation. Is this something that the version of belief holism under consideration can make sense of? It is hard to see that it can. The only difference between W and W^* , we are supposing, is the absence of S 's belief that it rained in New York on Jan. 1, 1900: W is just like W^* with this belief 'added on'. How does that belief become part of S 's single overall representation in W ? Why doesn't S in W have a single representation that is identical to the single overall representation she has in W^* , plus, in addition, the belief that it rained in New York on Jan. 1, 1900? By virtue of what does this belief, instead, enter into composition relations with S 's other W -world beliefs, such that she has in W a single, overall representation? This version of belief holism does not appear to say.

The variety of belief holism just considered is one on which the subject in question could lose any particular belief while still retaining the belief that a black cat is on a mat. One might think this is too permissive. Could she lose the belief that something exists, or that

there is at least one cat? It is plausible to think that she could not, that one could not believe that a black cat is on a mat without believing that something exists or that there is at least one cat. Consider, then, a different version of belief holism, one that says not merely that a subject has many beliefs if she has any beliefs, but also that there are certain specific beliefs a subject must have if she is to believe certain things. Let us grant, for instance, that a subject must believe that there is at least one cat if she believes that a black cat is on a mat. Does this different version of belief holism explain a subject's having one overall representation?

It seems to me that this version runs into the same worry as the first. In world W S believes a black cat is on a mat, that at least one cat exists, that it rained in New York on Jan. 1, 1900, and numerous other things $p_1 \dots p_n$. In W^* , S believes everything she believes in W except she does not believe that it rained in New York on Jan. 1, 1900. W is just like W^* with this belief 'added on'. If S 's belief system in W^* is or composes one representation, how does the belief that it rained in New York on Jan. 1, 1900 get to join or become a part of this representation in W ? Why doesn't S in W just have one (almost overall) representation that is identical to the overall representation she had in W^* , and then, in addition, the belief that it rained in New York on Jan. 1, 1900? Why does this belief about New York join, or become part of, her overall representation, instead of standing apart? This second version of belief holism seems, like the first, to have nothing to say on the question.

The two versions of belief holism just considered are ones on which it is intelligible to imagine a subject losing or acquiring a single belief, the belief that it rained in New York on Jan. 1, 1900. Although we commonly speak in this way, there is good reason to think this is impossible. It is reasonable to think that every world in which I believe that it rained in New York on Jan. 1, 1900 differs from every world in which I fail to believe that it rained in New York on Jan. 1, 1900 by more than just the presence or absence of this particular belief.

We do not need to settle this question: let us just grant, for the sake of argument, a third version of belief holism. This one says not merely that one has many beliefs if one has any beliefs, nor merely that there are some specific beliefs a subject must have if she is to believe, for instance, that a black cat is on a mat. On this version, every belief had by a subject is such that it can only be had by the subject if the subject has various other specific beliefs. On this third version of belief holism, therefore, worlds like W^* described above are impossible. That world was introduced via appeal to W , a world in which a subject believed many things, among which was that it rained in New York on Jan. 1, 1900. World W^* was defined as a world just like W , minus S 's belief that it rained in New York on Jan. 1, 1900. The third version of belief holism says this world is impossible, for if S lost that belief she would lose other beliefs as well. Does this version of belief holism do any better a job of explaining a subject's having a single, overall representation?

It seems to me that it does not. For it seems that there has to be some world analogous to the original W^* , a world suited to play the same role in the dialectic as W^* does. Call W^{**} that world in which S believes everything she believes in W except (i) that it rained in New York on Jan. 1, 1900, and (ii) whatever else she could only believe if she believed that it rained in New York on Jan. 1, 1900. Once again, W is just like W^{**} but with some set of beliefs added on. We again can ask: Why do S 's beliefs in W form a single, overall representation, rather than being simply such that S has one representation identical to the representation she has in W^{**} , plus, in addition, some other beliefs (those beliefs she has in W but not in W^{**})? How do those beliefs get to be part of her overall representation in W ? Belief holism explains why S cannot have merely one belief, but it does not seem to have anything to say about why it is that a subject's beliefs form a single, overall representation. Or, to put the point differently, belief holism does not seem to entail that a subject's beliefs

form a single, overall representation; if we want a proof of this, or reason to believe it, we must look elsewhere.

There is a further point to be made concerning the potential for belief holism to provide an account of global representations. In the earlier discussion describing the view, belief holism was interpreted as the claim that there is a kind of metaphysically necessary connection between beliefs or belief states. A belief such as that one is seeing a black cat could simply not be had in the absence of sufficiently many other beliefs about cats, blackness, seeing, and so on. However, on closer look this does not license the conclusion that *any possible belief* can only be had if certain other beliefs are had. The reason for this is that the mere fact that concepts such as ‘cat’, ‘black’ and ‘on-ness’ are connected to other concepts does not show that all possible concepts are. In other words, Davidson has not given us reason to believe that to have *any* belief requires that one have many beliefs: rather, he has given us reason to believe that to believe that one is seeing a black snake, and to have beliefs like that in relevant ways, requires that one have many other beliefs. That is, he has not demonstrated, so far as I can tell, that *every* possible belief is like the belief that one is seeing a black snake in the relevant regard.

Let me explain what I mean. It may be that the concepts we have are interconnected in various ways because the entities that fall under them or the properties they pick out, so to speak, are interconnected in various ways. Consider Davidson’s claim that one cannot wonder whether one is seeing a snake without having many true beliefs about snakes, such as that they lack legs, that they are long and skinny, that they move a certain way, etc. It may be that this dependency holds because the property of being a snake is instantiated only if these properties are instantiated. That is, it is part of the nature of snakes that the property of being a snake is instantiated only if many other properties are instantiated – properties

such as ‘lacking legs’, ‘being long and skinny’, ‘moving a certain way’, etc. It is not clear that all possible entities or properties are like this. Perhaps there is a possible world containing entities and properties that are not interconnected in this way. What we may need to imagine is a world containing a number of simple (atomic) entities instantiating certain simple (atomic) properties, where the properties in question do not depend for their instantiation on the instantiation of a great number of other (positive) properties. Suppose some simple entity instantiates G, the property of being a certain way, where the identity conditions of being G do not involve appeal to other properties. It is not clear that such a world is possible – it is not clear, that is, that it is possible for *properties* to be independent of each other in this way. But is equally not clear that it is impossible, which is what has to be true, it seems to me, if belief holism is going to be a claim that holds of all believing, actual or possible, which is what I read Davidson as arguing. (Alternatively, it could be that such properties are possible but it is metaphysically impossible for any believer to form beliefs about them. This is an interesting epistemic possibility. If Davidson is right that belief is holistic, but such properties are possible, then he may be committed to the view that they are beyond the reach of belief.) Whether such properties are metaphysically possible would seem rather to be an issue in ontology that warrants further study. If they are possible, then it is not clear that belief holism would have to be true in that world – or, at least, it is not clear that to have beliefs about such entities one would have to have many other beliefs. This is relevant to the question of whether belief holism explains or accounts for a subject’s having a single overall representation because it is reasonable to think it would be at least metaphysically possible to have such a representation in a world as described. I cannot discuss this in more detail here, other than to note that the positive solution I propose later

in the dissertation would account for a subject's having a single, overall representation in such a world.

Holism of Belief Content

A second variety of holism that Davidson endorses is holism of belief content, according to which, roughly, the content of a belief state (or of any propositional attitude state) “depends on its place in the whole network” (2004a, p. 15). It is not just that I cannot believe a cat is on the mat without believing many other things about cats, about mats, and about the relation of on-ness. It is that the content of my belief that a cat is on the mat is in some sense determined by what I believe about cats, mats, and on-ness (and, perhaps, determined by everything else that I believe, to greater or lesser degree). What I believe about cats, mats, and on-ness is, in turn, determined in some sense by my other beliefs as well. The content of each is in some sense determined by the content of all.

It is not possible to illustrate these complex mutual determination relations in detail, since there are, on this picture, far too many, but perhaps it is possible to show something of what Davidson has in mind. My belief that a cat is on the mat has the content it does in virtue of, or at least in concert with, my many other beliefs' having the content that they do, especially those concerning cats, mats, and on-ness. One thing I believe about cats is that they are a product of evolution. I have many beliefs about evolution, of course; picking just one, I believe that some animals prospered because genetics disposed them to be more fit to the environment of their birth and environmental conditions remained stable, whereas other animals prospered because genetics disposed them to be more fit to deal with some environmental change that happened after their birth. This belief involves such concepts as 'environment', 'birth', and 'fitness', about which I have many more beliefs, and so on. My

belief that a cat is on the mat has, in some sense, a different content than it would if I did not believe cats were products of evolution, or believed they were but thought evolution worked differently than I take it to work, and so on.

We can perhaps get a better sense of what Davidson is getting at by looking briefly at some of what lies behind it.⁶ In claiming that the content of a subject's beliefs is in some way mutually determined, Davidson echoes, or builds on, earlier work by Duhem, Hempel and Quine. Hempel had posited a kind of holism of content in order to explain or account for *confirmation holism*, first defended by Duhem. Confirmation holism holds that a theoretical sentence of science entails certain observation sentences only in concert with many other theoretical sentences. Some theoretical sentence of biology, for instance, only entails that a formation of a certain kind will appear in the microscope if one holds fixed, or assumes, numerous other theoretical sentences – that, for instance, the microscope functions a certain way, that light behaves a certain way, that one's brain will function as it normally does, and so on. (Suppose one sets out to test a hypothesis, say that adding a certain pigment P to a cell will cause it to behave in way X. If one adds the pigment and the observation is not of the cell behaving as expected, this could be because the theoretical sentence "If pigment P is added to a cell, it will behave in way X" is false. But it could just as well be that this sentence is true and that various other theoretical sentences, such as "Adding pigment P to a cell will

⁶ I should note that Davidson's motivations for endorsing holism of belief content are complex, and I only gesture at the history of the idea here. Additional support is found within Davidson's theory of interpretation, which I say a bit more about below, as well as within the systematicity of language. On this last point, a quotation that illustrates how holism of belief content rests on the relation between word and sentence is the following: "If sentences depend for their meaning on their structure, and we understand the meaning of each item in the structure only as an abstraction from the totality of sentences in which it features, then we can give the meaning of any sentence (or word) only by giving the meaning of every sentence (and word) in the language. Frege said that only in the context of a sentence does a word have meaning; in the same vein he might have added that only in the context of the language does a sentence (and therefore a word) have meaning." (Davidson, 2001c, p. 22)

not make my eyes deceive me as to whether the cell is behaving in way X” are false.) The point was originally put by Pierre Duhem this way:

A physicist decides to demonstrate the inaccuracy of a proposition; in order to deduce from this proposition the prediction of a phenomenon and institute the experiment which is to show whether the phenomenon is or is not produced, in order to interpret the results of this experiment and establish that the predicted phenomenon is not produced, he does not confine himself to making use of the proposition in question; he makes use also of a whole group of theories accepted by him as beyond dispute. The prediction of the phenomenon...does not derive from the proposition challenged if taken by itself, but from the proposition at issue joined to that whole group of theories: if the predicted phenomenon is not produced, not only is the proposition questioned at fault, but so is the whole theoretical scaffolding used by the physicist. The only thing the experiment teaches us is that among the propositions used to predict the phenomenon...there is at least one error; but where this error lies is just what it does not tell us.” (Duhem 1954, p. 185, quoted on p. 232 in Achinstein).

Hempel, who along with Quine was responsible for popularizing Duhem's thesis in Anglo-American philosophy,⁷ sought to account for this kind of holism by appeal to a kind of meaning holism according to which there are semantic interdependencies between statements.⁸ Quine also endorsed confirmation holism (claiming that “the unit of empirical significance is the whole of science” (1951, p. 42)), and took it to entail something like content holism, which found expression in his attack on the analytic/synthetic distinction. The best short summary of this view that I know of is one offered by Neil Wilson:

We may gloss Quine’s thesis about the analytic-synthetic distinction as: There is no sharp line between what properly belongs in a dictionary and what properly belongs in an encyclopedia. (Wilson 1967, p. 63. Wilson attributes the phrase to Jerrold Katz.)

⁷ Although confirmation holism is perhaps best known through the work of Hempel and Quine, both explicitly acknowledge Duhem's early formulation of the view. As Hempel put it, “It is remarkable that already in 1906 Pierre Duhem, in *La Théorie Physique*, presented the basic idea in a clear way, justified it thoroughly, and investigated its consequences”. (2000, pp. 281-282)

⁸ “In order to understand “the meaning” of a hypothesis within an empiricist language, we have to know not merely what observation sentences it entails alone or in conjunction with subsidiary hypotheses, but also what other, non-observational, empirical sentences are entailed by it, and for what other hypotheses the given one would be confirmatory or disconfirmatory. In other words, the cognitive meaning of a statement in an empirical language is reflected in the totality of its logical relationships to all other statements in that language, and not to the observation sentences alone.” (Hempel 1950, p. 59)

On the traditional view, dictionaries define words, whereas encyclopedias offer empirical information about the things these words refer to. Davidson builds on Quine's insistence that there is no sharp distinction between what is properly part of the meaning of a term and what is properly part of what we know about the entities that fall under the term by claiming, roughly, that the content of any of a subject's belief states is given by the totality of the content of the subject's belief states. What a subject believes when she believes, for instance, that a cat is on the mat is given in part by everything that she believes about cats.

There is something powerful and compelling about Davidson's holism of belief content. There is some force to the idea that what I mean by "cat" is something like that entity that is furry and lithe, that stalks mice, that has a certain kind of evolutionary history, and so on, where what I mean by "furry" or "lithe", or "stalks mice", or "evolution" is determined in turn by the beliefs I have about furriness, liteness, mouse stalking, evolution, and so on. Once we move beyond this merely impressionistic story, however, it is quite unclear what exactly holism of belief content amounts to.

One problem is understanding the metaphors by which the position is usually described. Peter Pagin notes as much, and indicates some of what is puzzling, in a discussion of meaning holism (of which holism of belief content is a variant):

MH is sometimes characterized by saying that the meaning of a sentence, or a belief state, is given by its place in a 'network', 'web', 'pattern', 'space' or 'system' of sentences or beliefs. The network metaphor is not cashed in, however, but only used as an illustration. The illustration is somewhat misleading, since it suggests that—like the nodes in a network are pairwise connected with lines—the relevant interconnections between expressions consist in a large number of *binary* relations, and also that a metric of *distance* between the nodes can be defined on that basis, so that we have a well-defined notion of a *location* in the network (directly connected expressions are supposed to be 'closer' than indirectly connected ones). By contrast, the relations actually considered in theories of meaning determination are more complex. (2006, p. 221)

A second problem, and in my view an equally serious one, is understanding the relevant notion of *determined*. In saying that the content of a belief is determined by its relation to other beliefs, Davidson has more in mind than merely that one could not believe that G is F unless one had enough beliefs about G and about F-ness. Holism of belief content is not merely the claim that I can have a certain belief only if I have sufficiently many other beliefs. For this is just the claim of belief holism, the thesis considered above, and Davidson takes belief holism to support holism of belief content, which implies that holism of belief content is something other than mere belief holism:

Such interdependence [of content, i.e., holism of belief content] is already supported by the arguments for a multitude of thoughts [i.e., belief holism]. Thus if to believe one is seeing a snake requires that one have many beliefs about the nature of snakes, then it follows that if enough of those further beliefs were to change, so would the belief that one is seeing a snake. (2004a, p. 14)

We can also see that belief holism is distinct from holism of belief content by noticing that the latter entails the former, but not vice versa. If the content of a subject's beliefs is determined by her other beliefs, then if she has any contentful belief state (i.e., if she has any belief at all) she has many beliefs. But the reverse entailment fails: the fact that a subject necessarily has many beliefs if she has any does not entail that the content of her beliefs is mutually determined, for it could be that belief holism is true for epistemic rather than semantic reasons.⁹

⁹ Pagin notes just this point, although he speaks of dependence rather than of determination: "In the case of belief holism, the stress is on the conditions for a person to have a belief with such and such a content. This is clearly a different question from that concerning how belief states depend on each other for having their content fixed. If there is a dependence, so that one belief state cannot have a particular content unless it is somehow connected to other belief states with appropriately related contents, then belief holistic claims follow. But the converse doesn't hold. There can be other reasons for belief holism than MH [meaning holism, i.e., holism of belief content]. For instance, it is reasonable to claim that to have a belief that a gun is loaded the believer must minimally be able to distinguish guns from other things. This may then be combined with two further claims: first, that the only way of having that ability is having a grasp of functional features of guns (as distinct from perceptual features), and second, that grasp of functional features requires further beliefs. The need for having further beliefs is then *epistemological* rather than *semantic*: other beliefs are needed for some particular cognitive capacity. If the dependence of some beliefs on others isn't semantic in nature, we don't have an example of MH." (2006, pp. 215-216)

The passage just quoted also captures a third thing that is ambiguous or puzzling about Davidson's proposal, for it is unclear what Davidson means when he talks about "beliefs changing". What would it be for my beliefs about the nature of snakes to change? Suppose that I believe at t , some time at which I report that I believe I am seeing a snake, lots of things about the nature of snakes, and that at some later time t' don't have those beliefs but have instead some other beliefs (leaving open the question of whether those later beliefs are about snakes or are about other creatures similar in some ways to snakes). Either at t' my new beliefs are about snakes or they are not. If they are in fact about snakes (if the total changes to my belief system haven't been great enough to make it the case that I am not thinking about snakes), then it seems that when I say, at t' , "I am seeing a snake", I am reporting that I believe that I am seeing a snake. In this case, it would seem that my belief has not changed – I believed at t that I am seeing a snake, and I believe this at t' as well. Suppose now, instead, that the total changes to my belief system make it the case that my new beliefs are not about snakes, but about some other creatures similar in certain ways to snakes. In this case it doesn't seem that my belief that I am seeing a snake has changed. It has, rather, gone out of existence. I no longer believe that I am seeing a snake; I believe, instead, that I am seeing a snake* (which is some creature similar in certain ways to a snake). For this reason, it is unclear why Davidson thinks that belief holism leads to holism of belief content. We can just as well imagine that when the 'supporting' beliefs change, what changes is not the content of the belief that one is seeing a snake, but rather whether one has that belief at all.

There is much more than can be said about holism of belief content¹⁰ – the topic is far beyond the scope of a dissertation chapter. My goal, however, is not to gain a complete understanding of this Davidsonian claim, nor to engage in Davidson scholarship. Rather, I want to see whether there is any connection between holism of belief content and global representation.

Prima facie, there might seem to be a close connection. Quine's holism is often interpreted to be one on which what has 'empirical content' is a whole theory, with individual statements within it having content only derivatively (Pagin 2006, p. 3). We might, along the same vein, think that holism of belief content entails that what has content, primarily, is a subject's entire belief system, and that her individual belief states have content only derivatively. Moreover, holism of belief content may seem to do a good job of explaining the 'picture' metaphor through which the claim that a subject has a global representation of the world finds natural expression. According to holism of belief content, one has any particular belief one has only by virtue of, or in concert with, one having the other beliefs that one has. Consider a painting of the Declaration of American Independence. It seems reasonable to think that, in the case of the painting, individual splotches of paint have the content they do (i.e., represent what they do) only by virtue of or in concert with their occurring within the whole painting. (Think of some one-square millimeter splotch of blue paint, which represents some part of Thomas Jefferson's jacket. It is reasonable to think that this splotch only represents some part of his jacket by virtue of or in concert with all the other splotches of paint representing what they do.) In some sense, the splotches all hang together, representing the event of the signing of the American

¹⁰ And much more that has been said: important discussions of semantic holism include Fodor and Lepore 1992; Fodor and Lepore 1993; Bilgrami 1998; Guttenplan 1994; Lormand 1996; Block 1995; Devitt 1993; and Devitt 1995.

Declaration of Independence, and each little splotch represents what it does in some sense because the whole represents what it does. (I am here not making a definitive claim concerning the metaphysics of representation in paintings. Rather, the idea is only to illustrate why one might think that holism of belief content is connected in an important way with the unity of belief.) One might think that this is just what is going on with belief, if Davidson is right that the content of each belief is determined by the content of the whole.

There is something important and interesting about this analogy, but I do not think that mutual determination of content is what explains why or entails that a subject has an overall representation. How does the fact (assuming, for the sake of argument, that it is a fact) that the contents of my many beliefs are mutually determined show or entail that I have one representation that has as part of its content the content of my many beliefs? We can put pressure on this entailment by considering the following thought experiment: There are two books designed or written in such a way that their content is mutually determined. Perhaps, for instance, they refer to each other in a certain way. Book A includes, for instance, the sentence “Looking over his left shoulder, S (the first person to speak in chapter five of Book B) scurried quickly down the street and ducked into an open doorway.” Here the idea is that this sentence within Book A is about S, a person who appears in Book B. S’s identity conditions, so to speak, are in part given by Book B – the first person to speak in chapter five of Book B is Theodore, a lean, angry man who chain-smokes Marlboro Lights, etc. We can suppose the converse is true as well: Book B says that Theodore is a lean, angry man who chain smokes Marlboro Lights, and that he is a lean, angry chain smoker because of the trauma endured on page 11 of Book A. The content of Book A and Book B is mutually determined because Book A picks out Theodore in part by appeal to the content of

Book B, and the Theodore of Book B has some of the properties he does (or, to put it another way, is the person he is) in part by appeal to the content of Book A.

Is there in this case anything that is a representation that represents everything that Book A represents and everything that Book B represents? There does not seem to be. Book A represents a complex of states of affairs, as does Book B. Book A represents the complex of states of affairs it does partly in virtue of Book B's representing the complex of states of affairs that it does, and vice versa. But this does not, it seems to me, entail that there is anything that represents all of what Book A and Book B represents. If this is right, we have mutual determination (to some degree, at least) of semantic content without there being a single, overall representation.

Is it possible that, contrary to what has been said, Book A and Book B are somehow, together, a single overall representation? Certainly it seems possible that they could be. Perhaps they have a single maker, an artist who intends them in some sense to be a single work, as we might image a trilogy or a triptych is. However, if holism of belief content is to explain a subject's having a single, overall representation, then mutual determination of content must be sufficient for their being a single, overall representation. In other words, while there could be cases where books like Book A and Book B together form a single representation, if holism of belief content is to do the job we need it to do, then it must be that every pair of books like Book A and Book B together form a single, overall representation. It does not seem to me that this is the case. I should also note that it will not help to suppose that a subject who reads and remembers the details of both books will have a single representation that has as its content all the content of Books A and B. For in this

case the unity belongs to belief, and not to the books themselves in virtue of their content being mutually determined.¹¹

There is perhaps one way that a proponent of holism of belief content could argue that such holism entails that a subject has, or explains a subject's having, a global representation of the world. It requires adopting an implausibly strong version of holism of belief content, however. I will sketch the argument, but not develop it in any detail since (i) the strong version of holism on which it rests is implausible, and (ii) Davidson explicitly rejects this version of holism as well.

Suppose the content of a belief is maximally fragile or instable: if B is a belief with content *p*, there is exactly one belief system to which B can belong. That is, no two believers agree on anything unless they agree on everything. This version of holism of belief content has been widely scorned,¹² and Davidson himself is careful to distance himself from it.¹³ Yet this version of holism of belief content, if it were true, could plausibly be taken to hold out some promise for explaining the unity of belief. The first thing to note is that it, together

¹¹ I should also note that in the case of the painting, what makes the painting one representation is not the mutual determination of content. If anything, that the content of various splotches of paint is mutually determination rests on, rather than explains, the painting's being one representation.

¹² Most notably, perhaps, by Jerry Fodor – see especially Fodor and Lepore 1992. It is unclear how many philosophers sympathetic to semantic holism endorse this strong version of holism. Michael Devitt argues that most do, at least implicitly, whereas Jones, Mulaire and Stich (as Devitt notes) argue that Fodor's attack on this version of holism amounts to an attack on a Straw Holist. See Devitt 1993 and Jones, Mulaire, Stich 1991. See also Jackman 1999 for a discussion of the relation between the instability thesis and semantic holism.

¹³ Davidson often says things that make it sound as if he endorses this strong version of holism of belief content. E.g. "[A] proposition cannot maintain its identity while losing its relation to other propositions". (2004f, p. 204). (Here Davidson speaks of the identity of propositions, but the context shows that his point is rather about belief.) However, he explicitly makes clear that he means only that the position a belief occupies within a network is determined what other beliefs one has, but that a belief can occupy the same position even if some of the beliefs change. How this is possible, on his view, is described in the following passage:

Here is an analogy: any *one* change in the tension of one part of a spider's web will change the position of many parts of the web (all, in fact, except the anchor points). But given an exchange, many possible adjustments in the tension elsewhere would preserve the position of most parts. Or consider adjustments in the center of gravity of an airplane. If some person changes his seat in an airplane, the relation of every object in the plane and every part of the plane to the center of gravity changes. But a single compensatory move will restore the center of gravity, and hence the relations of all objects to the center. (2004a, p. 14)

with an attractive metaphysical principle, yields the view that a subject's beliefs at a time are not wholly distinct. The plausible principle is Hume's dictum to the effect that there are no metaphysically necessary connections between distinct existences.¹⁴ Suppose the variety of holism of belief content on which belief content is maximally fragile is true. This entails that if $B_1 \dots B_n$ are beliefs had by some subject S , there is no world in which any of $B_1 \dots B_n$ are instantiated in the absence of all of $B_1 \dots B_n$. Hume's dictum would seem to entail that such beliefs fail to be wholly distinct. Perhaps one could argue by appeal to the Humean dictum that $B_1 \dots B_n$ constitute *one thing*, and, further, that that thing is representational, in that it represents exactly whatever is the case if the beliefs $B_1 \dots B_n$ are true. This might then yield the conclusion that believing subjects necessarily have a global representation. I will not pursue this argument further, however, since I think the version of holism of belief content on which it rests is false. Moreover, my goal in this chapter is principally to consider whether Davidson's well-known holism of the mental involves or explains a subject's having a global representation of the world. Since Davidson explicitly denies that belief content is radically fragile in the way it would need to be were this 'Humean' argument to work, he could not endorse this argument. In other words, we need not determine whether the conjunction of Hume's dictum and radical fragility of belief content entails that a subject has a global representation of the world in order to determine whether Davidson's holism of belief content entails that a subject has such a representation.

Interpretive Holism

A third thesis of Davidson's is interpretive holism, according to which we cannot interpret the actions and utterances of an agent without attributing to the agent a full system of

¹⁴ "There is no object, which implies the existence of any other if we consider these objects in themselves." *A Treatise of Human Nature*, Book I, Part III, § V.

propositional attitudes. Although this version of holism seems to have the least to do with the notion of global representation, I want to consider it nonetheless. There are two reasons for this. First, it is in the context of discussing interpretive holism and the closely related principle of charity that Davidson makes pointed appeal to the concept of a global representation. And second, because although I do not think Davidson's discussion of global representation, in the context of interpretative holism, is ultimately satisfying, he appears to be sensitive to some of the problems I take to be most interesting.

Davidson's claim that interpretation is holistic is akin to the confirmation holism explored by Duhem, Hempel and Quine. Where they focus on what is required to confirm a theoretical sentence, Davidson focuses on what is required if we are to interpret the utterances (and thereby, the propositional attitudes) of a speaker. The basic idea is that in interpreting the utterances of others we have to posit an entire system of propositional attitudes. This might be in part for reasons discussed earlier – we could not attribute the belief that there is a cat on a mat to a speaker unless we also attributed to her numerous other beliefs about cats, mats, etc. But it is also because, Davidson claims, an agent's propositional attitudes are constrained by norms of rationality. We could not interpret an agent as believing that there is a cat on the mat and as believing that if there is a cat on the mat and she does not do anything to get the cat off the mat something terrible will happen without attributing to her other propositional attitudes and mental attributes. We might attribute to her a desire to get the cat off the mat, coupled with an intention to do so. Or we might attribute no intention, but explain this by attributing a desire for something terrible to happen, or by attributing a desire that something terrible not happen but a belief that efforts to remove the cat would be futile. The general idea is that attributions of beliefs, desires, other propositional attitudes and other mental attributes come as a package deal, with any

possible package constrained by broad norms of rationality. In a passage echoing Duhem, Hempel and Quine's discussions of confirmation holism within science, Davidson says:

[Interpretative] Holism is forced on us by the fact that what can fairly be treated as evidence for any one of the propositional attitude can count as evidence only if untested assumptions are made about the others....The key to the solution for simultaneously identifying the meanings, beliefs, and values of an agent is a policy of rational accommodation, or a principle that Quine and I, following Neil Wilson, have called in the past the principle of charity. This policy calls on us to fit our own propositions (or our own sentences) to the other person's words and attitudes in such a way as to render their speech and other behavior intelligible. This necessarily requires us to see others as much like ourselves in point of overall coherence and correctness." (2004b, p. 35)

Davidson's interpretative holism, an epistemological thesis or a thesis with the philosophy of language, has a metaphysical twin, which may be called inter-attitudinal holism – the claim that beliefs, desires, and other propositional attitudes can only be had together with appropriately complementing attitudes. It is not merely that interpretation demands that we attribute to others a full array of propositional attitudes constrained by norms of rationality, it is that they could not have a belief, desire, or other propositional attitude without having such an array of attitudes.

Does interpretive holism entail that a subject has a global representation of the world? Again, the answer would seem to be no. Interpretative holism and the closely allied inter-attitudinal holism claim that propositional attitudes (beliefs, desires, hopes, fears, etc.) come as a package deal, in the sense that a subject cannot be merely a believer, or a desirer, or a hoper, or a fearer, etc.: subjects necessarily have a range of propositional attitudes if they have any at all. As Davidson memorably puts it, "the intrinsically holistic character of the propositional attitudes makes the distinction between having any and having none dramatic" (1982, p. 327). This does not seem to say anything about whether a subject has a global representation of the world. Moreover, if belief holism – the claim that a subject necessarily

has many beliefs if she has any beliefs – does not entail that a subject has an overall representation, then interpretive or inter-attitudinal holism – the claim that a subject necessarily has a vast, diverse array of propositional attitudes if she has any propositional attitudes – should not either. In both cases what is missing is any reason to think the necessity of a plurality (of beliefs, or of propositional attitudes generally) entails or explains the existence of a singleton (an overall representation or picture of the world).

It is worth noting, however, that it is in the context of interpretive holism that Davidson invokes the notion of a global representation of the world, for he argues that interpretation requires that we attribute to an agent we interpret “a picture of the world” that is largely true and largely in accord with ours. It is important here to quote at some length from a number of different passages, for what I am trying to do is show that Davidson explicitly appeals to the concept of a global representation in outlining and explaining interpretive holism and the principle of charity:

My main point is that our basic methodology for interpreting the words of others necessarily makes it the case that most of the time the simplest sentences which speakers hold true *are* true....Each speaker can do no better than make his system of beliefs coherent, adjusting the system as rationally as he can as new beliefs are thrust upon him. But there is no need to fear that these beliefs might just be a fairy tale. For the sentences that express the beliefs, and the beliefs themselves, are correctly understood to be about the public things and events that cause them, and so must be mainly veridical. Each individual knows this, since he knows the nature of speech and belief. This does not, of course, tell him *which* of his beliefs are true, but it does assure him that his overall picture of the world around him is like the picture other people have, and is in its large features correct. (Davidson 1986, p. 322)

Why must our language - any language - incorporate or depend upon a largely correct, shared, view of how things are? First consider why those who can understand one another's speech must share a view of the world, whether or not that view is correct. The reason is that we damage the intelligibility of our readings of the utterances of others when our method of reading puts others into what we take to be broad error. We can make sense of differences all right, but only against a background of shared belief. What is shared does not in general call for comment; it is too dull, trite, or familiar to stand notice. But without a vast common ground, there is no place for disputants to have their quarrel. (Davidson 2001a, pp. 294-295)

[O]ur view of the world is, in its plainest features, largely correct....What cannot be the case is that our general picture of the world and our place in it is mistaken, for it is this picture which informs the rest of our beliefs and makes them intelligible, whether they are true or false. (Davidson 1991, p. 160)

In sharing a language, in whatever sense this is required for communication, we share a picture of the world that must, in its large features, be true. (Davidson 1979, p. 294)

[T]he general outlines of our view of the world are correct; we individually and communally may go plenty wrong, but only on condition that in most large respects we are right. It follows that when we study what our language – any language – requires in the way of overall ontology, we are not just making a tour of our own picture of things: what we take there to be is pretty much what there is. (Davidson 2001e, xviii-xix)

There are two features common to these passages. The first is that interpreter and speaker must largely agree, where this is described or perhaps even cashed out as sharing a picture of the world. The second is that a subject's picture of the world is itself subject to veridical appraisal – that is, a subject's picture of the world can itself be true or mostly true or mainly true, where this may even involve the idea that the truth or representational fidelity of the picture itself comes in degrees. Both features are central to Davidson's philosophy of mind and language.

One might think that Davidson's use of the concept 'picture of the world', a concept I take to be largely the same as that of "global representation of the world", is merely incidental, in that he could just as well have said what he wished to say without appeal to this concept. To be sure, this line of objection or response would go, Davidson often speaks of interpretation as a matter of attributing to the interpreted subject a picture of the world that is largely true; this is the locution that makes it seem as if his theory of interpretation relies on the concept of a global representation. But in many other places he just speaks rather of interpreting an agent so that most of the subject's beliefs are true. For instance, again quoting at some length:

But of course it cannot be assumed that speakers never have false beliefs. Error is what gives belief its point. We can, however, take it as given that *most* beliefs are correct. The reason for this is that a belief is identified by its location in a pattern of beliefs; it is this pattern that determines the subject matter of the belief, what the belief is about. Before some object in, or aspect of, the world can become part of the subject matter of a belief (true or false) there must be endless true beliefs about the subject matter. (Davidson 2001d, p. 168, italics in original)

I urge that a correct understanding of the speech, beliefs, desires, intentions, and other propositional attitudes of a person leads to the conclusion that *most of a person's beliefs* must be true, and so there is a legitimate presumption that any one of them, if it coheres with most of the rest, is true....The agent has only to reflect on what a belief is to appreciate that *most of his basic beliefs* are true, and among his beliefs, those most securely held and that cohere with the main body of his beliefs are the most apt to be true. (Davidson 2001f, p. 146 & 153, italics mine)

It should now be clear what ensures that our view of the world is, in its plainest features, largely correct. The reason is that the stimuli that cause our most basic verbal responses also determine what those verbal responses mean, and the content of the beliefs that accompany them. The nature of correct interpretation guarantees both that *a large number of our simplest beliefs are true*, and that the nature of those beliefs is known to others. Of course many beliefs are given content by their relations to further beliefs, or are caused by misleading sensations; any particular belief or set of beliefs about the world around us may be false. What cannot be the case is that our general picture of the world and our place in it is mistaken, for it is this picture which informs the rest of our beliefs, whether they be true or false, and makes them intelligible. (Davidson 1991, p. 160, italics mine)

The natural inference to draw from these passages is that Davidson thinks that a subject has a picture of the world that is largely true if and only if most of her beliefs are true. The last quotation makes this particularly clear, as Davidson slides back and forth between talking of subject's picture of the world being true or correct and talking of a large number of her belief's being true. For this reason, one might think that Davidson is not *committed* to the claim that a subject has a global representation of the world – he is merely using the familiar metaphor as a figure of speech. On this view, in saying that interpretation requires that we attribute to an agent a picture of the world that is largely true, Davidson is saying nothing more than that interpretation requires that we interpret a subject such that most of her

beliefs or utterances are true. This line of response insists that we could eliminate the concept of a global representation from the claim Davidson is making without any loss.

I am not sure that this line of response is right. Davidson himself is not explicit on the issue. He clearly seems to think that a subject has a picture of the world that is largely true if and only if most of her beliefs are true. But he does not explain, as far as I can tell, whether he takes it that these locutions simply mean the same thing, or whether one's picture of the world and its being largely true reduces to one's having many beliefs most of which are true, or whether the reverse reduction holds. He does not even give us reason to think the biconditional is true. It is not obvious, in the sense of 'so evident as not to warrant discussion', that a subject's picture of the world is largely true if and only if most of her beliefs are true. Moreover, he often hedges the biconditional claim by saying that what is required is that most of one's simple or basic beliefs (as in the quotations immediately above) must be true. This qualification implies that one's picture of the world being largely correct is not merely a matter of the *number* of one's beliefs that are true, but that something else – some qualitative notion such as that of simplicity or centralness of beliefs – matters as well.¹⁵

What is especially important, however, is that Davidson elsewhere expresses doubt about the intelligibility of the notion of "most of one's beliefs being true", on the grounds that beliefs are in principle uncountable. This makes it very hard to see how he could think that one's picture of the world is largely true if and only if most of one's many beliefs are true. And it suggests that he may take a subject's having a global representation of the world

¹⁵ For discussion of a closely related issue, see Lepore (???) and Neale 2001, p. 187f. Lepore argues, and Neale accepts, that in interpreting an agent, an interpreter is not guided by a desire to minimize the *number* of false beliefs attributed. (Both accept that on Davidson's picture, numbers matter, but that "maximizing agreement" is not to be cashed out purely in terms of number.) It is worth noting that in the original formulation of the principle of charity, Neil Wilson put the principle in terms of numbers: "We translate in a way that makes the largest possible number of statements true." (Wilson 1959, quoted by Quine in *Word and Object* II.§13.]

to be in some important way more basic or fundamental than her having some number of beliefs. For instance, in a discussion of coherence, Davidson says:

So mere coherence, no matter how strongly coherence is plausibly defined, cannot guarantee that what is believed is so. All that a coherence theory can maintain is that most of the beliefs in a coherent total set of beliefs is true.

This way of stating the position can at best be taken as a hint, since there is no useful way to count beliefs, and so no clear meaning to the idea that most of a person's beliefs are true. A somewhat better way to put the point is to say there is a presumption in favor of the truth of a belief that coheres with a significant mass of belief. (Davidson 2001f, p. 138-139)

This is a very interesting passage, for Davidson expresses doubts about the intelligibility of the idea of counting beliefs, and gestures toward a solution by recasting the plural ("most of the beliefs in a total set") into the singular ("a significant mass of belief"). One possibility is that Davidson thinks that the idea that a subject has *many* beliefs is unintelligible, but the idea that a subject has a mass of belief, a kind of singleton, is not. He also appears to be sensitive to the fact that it is quite unclear what it is for most of one's beliefs to be true.

In the passage just quoted, Davidson does not make clear what his worry is concerning counting beliefs. That is, he says that there is no useful way to count beliefs, but he does not make clear why he thinks this is so.¹⁶ In other places, however, he says a bit more to show what the worry is. First, from "Thought and Talk":

No simple theory can put a speaker and interpreter in perfect agreement, and so a workable theory must from time to time assume error on the part of one or the other. The basic methodological precept is, therefore, that a good theory of interpretation maximizes agreement. Or, given that sentences are infinite in number, and given further considerations to come, a better word might be *optimize*. (2001b, p. 169, italics in original)¹⁷

¹⁶ It is also unclear why Davidson says there is no *useful* way to count beliefs. Why not just say there are no accurate or correct ways to count beliefs?

¹⁷ The 'further considerations' Davidson alludes to are these: some disagreements are more destructive of understanding than others. As Kathrin Glüer has explained "The question of what constitutes more or less agreement is a good and hard one, however. Not only is it difficult to count beliefs because of their holistic nature. Moreover, it seems very plausible to say that we have a deeper or more essential disagreement with someone who seems to be denying that water is wet than with someone who does not exactly know what a quark is. The first disagreement also is more damaging to understanding that person's beliefs than the second; therefore, it should count more, so to speak. Because of considerations like these, Davidson suggests to replace

This passage presents Davidson's main worry with maximization – and thus with the idea of counting beliefs – as resting on the alleged fact that a subject's beliefs are infinite in number. The idea, I gather, is that once we recognize that a subject has an infinite number of beliefs (or that there are an infinite number of sentences to which a believer would give assent), it no longer makes sense to say that, for instance, Smith shares more beliefs with Jones than with Brown since she shares an infinite number with each.¹⁸

Davidson has another worry about the intelligibility of counting beliefs, however, which is somewhat harder to understand, and perhaps somewhat deeper. This worry, I think, gives additional reason to think that Davidson may take a subject's having a global representation of the world to be in some sense basic. The worry finds expression in a passage, and accompanying footnote, in "What Metaphors Means". The passage reads:

But in fact there is no limit to what a metaphor calls to our attention, and much of what we are caused to notice is not propositional in character. When we try to say what a metaphor 'means', we soon realize there is no end to what we want to mention. If someone draws his finger along a coastline on a map, or mentions the beauty and deftness of a line in a Picasso etching, how many things are drawn to your attention? You might list a great many, but you could not finish since the idea of finishing would have no clear application. How many facts or propositions are conveyed by a photograph? None, an infinity, or one great unstatable fact? Bad question. A picture is not worth a thousand words, or any other number. Words are the wrong currency to exchange for a picture. (Davidson 1978, pp. 46-47)

[sic] talk of maximizing by talk of optimizing agreement." (2001, p. 91) See also footnote 15, this chapter. Davidson does not say much to explain how optimization of agreement is supposed to work, or why – and how it is the case that – some beliefs "count more". This is unfortunate, since so much of his philosophy of language rests upon this notion. Davidson may be right that optimization is a better concept than maximization, but there is little theoretical gain in switching from talk of maximization to talk of optimization, since the latter notion is no clearer than the former. A surprising number of philosophers have been happy to adopt talk of optimization. Bruce Vermazen, for instance, in a footnote to a paper assessing Davidson's claim that massive error is unintelligible, points to Davidson's worry about how to understand the maximization of true beliefs and says "Quantitative expressions in this paper, e.g., 'most', 'preponderance', and 'sum', are to be understood as measures of optimization (or 'pessimization'), not of numbers of beliefs." (Vermazen, 1983, p. 71)

¹⁸ In "The Problem of Objectivity", Davidson also makes it sound as if the worry about counting beliefs is largely based on his view that a subject's beliefs are infinite in number: See Davidson 2004a, p. 13.

Here it sounds just as if Davidson is making a point about metaphors, and about such forms of representation as pictures or maps. The idea is quite obscure, but seems to be that metaphors, pictures and maps do not convey any number of facts or propositions for the reason that what they convey is different in kind than what can be expressed in ordinary, literal discourse. However, in a footnote to this passage (occurring at the end of the sentence “When we try to say what a metaphor ‘means’, we soon realize there is no end to what we want to mention”), Davidson says:

Stanley Cavell mentions the fact that most attempts at paraphrase end with ‘and so on’ and refers to Empson’s remark that metaphors are ‘pregnant’ [*Aesthetic Problems of Modern Philosophy*, 79]. But Cavell doesn’t explain the endlessness of paraphrase as I do, as can be learned from the fact that he thinks it distinguishes metaphor from some (‘but perhaps not all’) literal discourse. I hold that the endless character of what we call the paraphrase of a metaphor springs from the fact that it attempts to spell out what the metaphor makes us notice, and to this there is no clear end. I would say the same for any use of language. (Davidson 1978, pp. 46-47)

This footnote is remarkable, for, as I read it, Davidson claims that any use of language – including ordinary, literal discourse – is like metaphor, maps and painting in being such that there is no end to what an occurrence of it draws to our attention. This is very puzzling for if, as Davidson seems to claim in this footnote, all language is like metaphor and photographs, then sentences, even of literal discourse, are “not worth a thousand words, or any other number.” The upshot would be that ‘words are the wrong currency to exchange for sentences’, which may appear to be outright contradictory.

I think that Davidson is making a deep and interesting point here, although the clarity with which he expresses himself leaves something to be desired. Strictly speaking what he is saying is contradictory, for in the body of the essay he affirms that the meaning or content of metaphors, photographs, paintings and the like cannot be cashed out in literal sentences, on the grounds that metaphors, photographs and paintings are different in kind

from literal sentences, while in the footnote he claims that literal sentences are just like metaphors, photographs and paintings in being such that their meaning or content has an endless character, such that it cannot be cashed out in literal sentences. However, I think his point, somewhat more clearly expressed, is this: to understand the content of a word is to understand the content of a language, where the content of a language is like the content of a picture. Neither a picture, nor a language, decomposes into some number of semantic items. There is no limit to the number of sentences that one could say in an attempt to express the content of a picture, but this is not because the content of the picture is identical to, in the sense of decomposes into, the content of an infinite number of sentences. Similarly, I think Davidson is saying, there is no limit to the number of sentences that one could utter in an attempt to express the content of a language, but this is not because the content of a language is identical to, in the sense of decomposes into, the content of an infinite number of sentences. It is because a language is like a picture, or is a picture.

There is much obscurity here, to be sure. This is one area within Davidson's thought where I think obscurity reflects depth, and I can do no more here to improve the lucidity of his concern. The relevant point for present purposes is that Davidson doubts the intelligibility of counting beliefs, and this makes it quite unclear what we are to make of the biconditional claim that a subject's picture of the world is largely true if and only if most of her beliefs are true. Even if Davidson's worry is only that a subject's beliefs are infinite in number, that he has this worry is enough to show that he may take it that there is a certain force or primacy to a subject's having a global representation of the world.¹⁹ The second worry, which is obscure but seems to conclude that language is, or is like, a picture, in being

¹⁹ In their 2005, Lepore and Ludwig briefly discuss the question: "What does it mean to say that we have largely true beliefs, or that most of our beliefs are true, if we have an infinite number of beliefs?" See Lepore and Ludwig 2005, p. 324.

such that it does not decompose into any number – even an infinite number – of propositions, gives even more reason to think the notion of a global representation or picture is central to Davidson's thought. This line of thinking deserves careful consideration and a more thorough examination than I can give it here. (This essay is not an essay on Davidson, which is what it would have to be for this issue to receive sufficiently close examination.) Without doing this, however, I think it is fair to say there is reason to think Davidson's work involves the concept of a global representation of the world, and that he has important and interesting things to say about it.

Compartmentalization

Davidson's pointed appeal to the notion of a global representation of the world suggests that he is sensitive to some of the issues I take to be most interesting concerning the many/one structure of belief. There is one further aspect of Davidson's account of the mental that is relevant to the unity of belief. This is his endorsement of compartmentalization, the thesis that in many instances one's beliefs are in some sense compartmentalized, fragmented, held apart, or walled off from one another. Davidson invokes the notion of compartmentalization in the face of what he takes to be an important problem. On the one hand (according to Davidson), there is something seriously wrong with attributing to an agent inconsistent beliefs or inconsistent sets of beliefs, desires, intentions and other propositional attitudes.²⁰ But at the same time there are clear cases where an attribution of some kind of irrationality or inconsistency is required. Davidson uses the notion of compartmentalization as a way of preserving both rationality/consistency (within compartments) while making room for irrationality/inconsistency (between members of

²⁰ A belief cannot *contradict* a desire or an intention, of course, since the latter two are not assertoric. But a set of beliefs, desires and intentions can be inconsistent in the sense that given the beliefs and desires, the intentions are the wrong ones to have, or given the beliefs and intentions, the desires are the wrong ones to have, etc.

different compartments, or across compartments). The idea is that rationality and consistency are required intra-compartment but not inter-compartment.

There are two key points to this picture. The first is that straightforward attributions of inconsistency are problematic. As Davidson describes this point:

Synchronic inconsistency requires that all the beliefs, desires, intentions, and principles of the agent that create the inconsistency are present at once and are in some sense in operation – are live psychic forces. It is by no means easy to conceive how a single mind can be described in this way. (2004e, p. 197)

We need not consider in detail Davidson's reasons for finding synchronic inconsistency problematic. One central idea, however, is that, on his view, logical relationships between propositional attitudes are in some sense constitutive of their contents – to desire that p , for instance, is to be such that if one believes that Φ ing will bring out p , then one intends to Φ (subject to appropriate *ceteris paribus* constraints). A second, closely related idea is that one can attribute propositional attitudes only by, roughly speaking, “making sense” of a person's actions and behavior – and a set of attributions will not make sense if it is internally inconsistent in some way.²¹

The second key point to Davidson's invocation of compartmentalization is that, despite the near-unintelligibility of synchronic inconsistency, there are situations where

²¹ Davidson takes these two points to complement each other, and argues for, or at least asserts, both points in numerous places. For instance: “Sentences, or rather the attitudes they express, owe their content, that is their meaning, to two things: their relations to other sentences or attitudes, and their relations, direct or indirect, to the world through perception. It is therefore impossible [presumably Davidson means instead something like ‘unwise’ or ‘a bad idea’] for an interpreter wholly to disregard the logical relations among a speaker's sentences or attitudes. This is not a matter of an effort on the part of agents to be consistent; it is a matter of their speech and behavior having the meaning they do because of how they hang together. Unless there is sufficient coherence, it is impossible to assign propositional contents to their speech, beliefs, desires or intentions. The interpreter's assumption of a degree of rationality on the part of those she wishes to understand is thus no more than a condition of understanding them at all.” (2001g, p. 89). As another example: “If someone believes Tahiti is east of Honolulu, then she should believe Honolulu is west of Tahiti. For this very reason, if we are certain she believes Honolulu is west of Tahiti, it is probably a mistake to interpret something she says as showing she also believes Tahiti is west of Honolulu. It is probably a mistake, not because it is an empirical fact that people seldom hold contradictory views, but because beliefs (and other attitudes) are largely identified by their logical and other relations to each other; change the relations, and you change the identity of the thought. Simple, easy to grasp logical relations can't be widely or often offended by a thinker and the workings of that thinker's mind still be identified as thoughts.” (2004g, p. 97)

something close to inconsistency appears to be the best explanation. For example, it seems we could have evidence that a person believes that p and that q , believes that if both p and q , then r , but believes that $\sim r$. On the face of it, this would seem to be evidence of irrationality – we would be attributing propositional attitudes in such a way that the subject comes out as inconsistent. Or suppose, to broaden the inconsistency beyond cases of belief, that someone believes that a certain action would bring about an end, desires that end all things considered, but fails to intend to perform the action in question. In this case too, something like irrationality appears to be at work. Examples Davidson cites include “wishful thinking, acting contrary to one’s own best judgment, self-deception, [and] believing something that one holds to be discredited by the weight of the evidence” (2004d, p. 170). In all such cases, he thinks, there is “not the failure of someone else to believe or feel or do what we deem reasonable, but rather the failure, within a single person, of coherence or consistency in the pattern of beliefs, attitudes, emotions, intentions, and actions” (2004d, pp. 170-1).

How to reconcile points one and two? How can it be both that propositional attitude ascriptions demand internal coherency and that people have inconsistent propositional attitudes? It is at this juncture that Davidson appeals to compartmentalization:

Indeed, if we are to explain irrationality at all, it seems we must assume that the mind can be partitioned into quasi-independent structures....To constitute a structure of the required sort, a part of the mind must show a larger degree of consistency or rationality than is attributed to the whole....The idea is that if parts of the mind are to some degree independent, we can understand how they are able to harbor inconsistencies. (2004d, p. 181)

My purpose here is not to assess Davidson’s theory of compartmentalization, nor to propose a different solution to the problem he identifies.²² Rather, the point of the discussion is to

²² It should be noted that Davidson is not entirely happy with his theoretical introduction of compartmentalization, saying “There is no question but that the precept of unavoidable charity in interpretation is opposed to the partitioning of the mind. For the point of partitioning was to allow inconsistent or conflicting beliefs and desires and feelings to exist in the same mind, while the basic methodology of all interpretation tells us that inconsistency breeds unintelligibility”. (2004d, p. 184) For another defense of a Davidsonian model of partitioning, see Pears 1984, chapter five. For arguments against such models, see Mele

show that Davidson's talk of compartmentalization makes it unclear whether he is properly thought to endorse the claim that a subject has a global representation of the world.

Why think this? Well, let us first note that it is not exactly clear what compartmentalization amounts to. The compartment metaphor is poorly understood, and Davidson's other expressions – "walled off", "exiled", "isolated", "place[d] out of bounds", "territories" (2004f, pp. 211-212) – do little to make the picture clearer. Davidson recognizes this problem, saying in a footnote to the passage quoted above:

...my highly abstract account of the partitioning of the mind...[has] nothing to say about the number or nature of divisions of the mind, their permanence or aetiology. I am solely concerned to defend the idea of mental compartmentalization, and to argue that it is necessary if we are to explain a common form of irrationality. I should perhaps emphasize that phrases like 'partition of the mind', 'part of the mind', 'segment', etc. are misleading if they suggest that what belongs to one division of the mind cannot belong to another. The picture I want is of overlapping territories. (2004d, p. 181)

Although it is not clear what exactly compartmentalization amounts to, on the most reasonable interpretation of the theory Davidson is offering, if A and B are two different sets of compartmentalized beliefs had by some subject, then while the beliefs within each compartment might compose an overall representation, the subject does not have any representation such that it is composed of all the beliefs in A and B (exhaustively, or along with other beliefs not in either A or B). This is the most reasonable interpretation because compartmentalization is supposed to 'keep apart' inconsistent beliefs, and it is hard to see that beliefs that compose one representation could be thought to be 'kept apart' in an appropriate sense. (Davidson himself says "two obviously opposed beliefs could coexist only if they were somehow kept separate, not allowed to be *contemplated in a single glance*" (2004h, p. 220, italics mine). The upshot is that Davidson appears to be arguing that a subject's beliefs,

1987, chapters 6 and 10, and Heil 1989.

if they are in any way inconsistent, do not compose an overall representation. It is not clear to me how this can be reconciled with other parts of his project, in particular his claim that interpretation demands that we attribute to an agent a picture of the world that is largely true and largely in accord with ours. My suspicion is that Davidson is not recanting his endorsement of the claim that a subject has a global representation of the world, but is rather conflicted on the issue, or expressing two thoughts that are in conflict with one another.

There is a final point to make about compartmentalization and the unity of belief. That is that Davidson's account of compartmentalization does not explain why beliefs within one compartment are within one compartment. That is, he tells us that inconsistency is sufficient for compartmentalization – that, necessary, if A and B (beliefs, desires, intentions, etc.) are inconsistent then A and B are in different compartments. But does this mean that if A and B are consistent, then they *do* belong to the same compartment? Davidson is not explicit about this, but presumably this would be his view. But what is missing is an explanation of why consistency is sufficient for co-membership within a compartment. Suppose my beliefs are compartmentalized, due to my harboring some inconsistencies, where the idea is that of many overlapping territories. Now suppose I acquire some new belief, say that my friend Affan is getting married, that is inconsistent with none of my propositional attitudes. On Davidson's account, or at least on the interpretation we are presently considering, according to which consistency is sufficient for co-membership within a compartment, my belief that my friend Affan is getting married would be a member of every one of my propositional attitude compartments. But why is that? Why isn't it a member of some compartments and not of others? Davidson's answer would be that it is a member of every compartment because it is consistent with each compartment. But it is

hard to see the explanatory force of this. What is it about consistency that *entails* or *ensures* co-compartmentalization?

Conclusion

At first glance, Davidson's discussion of mental holism may appear to be closely connected to the unity of belief, and he appears to endorse the claim that a subject has a global representation of the world. Upon closer look, his discussion of holism, whatever its merits, does not offer an account of the unity of belief, and Davidson appears to hedge or qualify his endorsement of the global representation claim. There are two aspects of Davidson's work, however, that I wish to draw special attention to, as they will both feature in the positive account I offer later in the dissertation. First, Davidson is very likely right to insist that to believe a proposition requires, with metaphysical necessity, that one believe many other propositions.²³ This alone is not enough to explain the unity of belief, as I have argued above. However, it will feature in an important way in my positive account. Second, it is worth emphasizing that when Davidson endorses or avows a commitment to the view that belief is *one*, he does so by appeal to the metaphor of a picture of the world. As I have argued above, he does not offer an account of the relation between a subject's (single) picture and her (many) individual beliefs. Nonetheless, the picture metaphor (which is by no means unique to Davidson) is interesting because of what it suggests about the content of the single representation. An ordinary picture of, say, some physical goings-on in the world, represents the objects it depicts as standing in some spatial relations to one another. We see each of the objects, and we see how they are spatially related to one another. The representation of that (their standing in such and such relations to each other) is a representation that necessarily includes within it the individuals represented. The idea that a

²³ See however Fodor and Lepore 1992 for a sympathetic approach to the possibility of punctate minds.

subject's single overall representation is a judgment the content of which is that various states of affairs obtain and that they stand in various relations to one another will feature in my own account (although the relations they are taken to stand in are not spatial relations.) But that has to wait for now.

LEWIS AND STALNAKER

A belief...is a map of neighbouring space by which we steer. It remains such a map however much we complicate it or fill in details.

– Frank Ramsey, “General Propositions and Causality”

In the previous chapter, I argued that Davidson appeals to the notion of a global representation of the world in describing interpretive holism, but that as rich and stimulating as his work is, we do not find therein a clear account of the nature of such representations, nor compelling reason to believe a subject has such a representation. Moreover, in committing to the compartmentalization of belief, Davidson appears to hedge or withdraw his endorsement of the claim that a subject has a global representation of the world. In the current chapter, I look at two other philosophers whose work could reasonably be thought to be closely connected to the unity of belief: Robert Stalnaker and David Lewis. In *Inquiry*, Stalnaker proposes an account of belief on which a belief state is “a conception of the way the world is”, and “individual beliefs are properties of such a belief state.” (1994, p. 84 & 69) Stalnaker does not explicitly argue that a subject has an overall representation of the world, but he takes it that there are something like overall belief states that in some sense contain or subsume individual beliefs, or have individual beliefs as properties. Lewis, on the other hand, defends, or at least expresses sympathy with, a view I call ‘belief monism’ in two distinct ways. First, in the course of defending modal realism, Lewis describes the semantics of belief he shares with Stalnaker as one on which belief is “singular, not plural”. (1986, p. 32) The

account occurs in the first chapter of *On the Plurality of Worlds*, as part of Lewis's attempt to catalogue the theoretical benefits of modal realism. The idea is roughly that the content of a person's system of belief is given by a class of possible worlds; the claim that belief is one rather than many falls out from the claim that what is characterized is an entire system of belief, considered as one rather than as many. On the face of it, this view is one on which belief is radically monistic. That is, Lewis's claim is not that a subject has many beliefs which in some sense compose an overall representation, or even that one has an overall representation which subsumes, but is in some sense prior to, one's many beliefs. Rather, the picture is one on which belief is merely one, rather than both one and many. This thesis appears, on the face of it, to be very close to the view, or to be a stronger version of the view, that in believing a subject has one overall representation, for it holds that a subject does not have many beliefs but rather something like one belief, one taking of the world to be some particular way. Second, Lewis elsewhere endorses, or at least expresses sympathy with, the thesis that belief is 'map-like', and takes this to entail that there are no such things as individual beliefs. Lewis neatly summarizes the view:

If our beliefs are 'a map...by which we steer', as Ramsey said, then they are to that extent not language-like. And to that extent, also, it is misleading to speak in the plural of beliefs. What is one belief? No snippet of a map is big enough that, determinately, something is true according to it, and also small enough that, determinately, nothing is true according to any smaller part of it. If mental representation is map-like (let alone if it is hologram-like) then 'beliefs' is a bogus plural. You have beliefs the way you have the blues, or the mumps, or the shivers. (1999, p. 311, italic in original)

This view, like the first, is one on which a subject's believing is one whole, a whole without any correct or natural decomposition into individual beliefs. I will discuss these two views of Lewis's separately, as they are importantly different. For ease of reference, I will call the first

‘the modal semantic theory’ and follow standard practice in calling the second ‘the map-like theory’.

The chapter focuses on two questions: (i) is Lewis or Stalnaker in fact offering a view of belief on which a subject has a global representation of the world?; and (ii) does either Lewis or Stalnaker provide an adequate theoretical account of what it is for a subject to have a global representation of the world? I will focus for the most part on Lewis, since he, unlike Stalnaker, explicitly claims that belief is singular rather than plural.

The balance of the chapter is structured as follows. Following a terminological point, I begin by outlining, in very general terms, Stalnaker’s possible world semantics of belief, which Lewis’s modal semantic theory builds on. I then turn to Lewis, and make clear why one might think his is an account on which a subject has a global representation of the world. Following this, I argue Lewis’s modal semantic theory of belief is not one on which, in most instances, a subject has a single, overall representation – for reasons that will be familiar from the chapter on Davidson – nor one that provides an adequate theoretical account of a subject’s having such a representation even in those cases in which, on the view, a subject does. Following this, I briefly return to Stalnaker, making similar points about his view that a “belief state” is “a conception of the world”. The chapter concludes with a brief discussion of Lewis’s map-like theory.

Belief Monism

I call the general picture of belief offered by Lewis ‘belief monism’ because of the analog with metaphysical monism, a thesis concerning the ontology of the world.¹ The relation between monism vs. pluralism as a metaphysical dispute and monism vs. pluralism as a

¹ Monism can be understood as a thesis about either objects or states of affairs, and as a claim about either existence or ontological priority. For sympathetic discussion of monism as a thesis about objects and ontological priority, see Schaffer 2007; Schaffer forthcoming a; Schaffer forthcoming b; Schaffer in progress.

dispute about the nature of belief is interesting and, to my mind, underexplored, especially given the renewed interest in metaphysical monism. We can distinguish four distinct metaphysical theses, each of which has a parallel in belief: (i) Nihilist atomic pluralism takes it that there are exactly as many objects as there are simples. This parallels a view about belief on which a subject's beliefs are merely many – that is, a view on which one has many beliefs but no overall representation composed out of simpler beliefs; (ii) compositional atomic pluralism takes it that there exists some number of simples, plus some number of composite objects composed out of these simples, and that the world is a composite object, the whole that is not a part of anything else. This parallels what I call the unity of belief thesis, according to which one has many beliefs and, in addition, an overall representation these many beliefs compose; (iii) Priority monism takes it that the world is ontologically basic, that the parts of the world are ontological dependent on the world as a whole. This parallels a view of belief on which one's overall representation is fundamental and one's many individual beliefs depend on or are derivative of one's overall representation; (iv) Existence monism takes it that there is just one thing – the world. The parallel with belief is any view on which belief is “singular rather than plural”, or one that holds that a subject has one picture of the world but not, properly speaking, many individual beliefs. Lewis's discussion of belief presents his theory as being a kind of existence monism concerning belief, whereas Stalnaker seems to have something akin to priority monism in mind.

Stalnaker's Semantics

Stalnaker's theory of belief, and of intentionality in general, is one that understands representational content by appeal to possible worlds. Such models began with Jaakko Hintikka's *Knowledge and Belief* (Hintikka 1962), but Stalnaker's version is the best known.

Perhaps the best way to briefly outline Stalnaker's possible world semantics of belief is to look at what he takes its basic motivation to be and at how he contrasts it with what he calls the linguistic picture of belief. Although Stalnaker believes a possible worlds model has various theoretical advantages, what drives it is a conception of belief as pragmatic. The idea is that belief is essentially connected to action, and action is essentially connected with drawing distinctions between possible (future) states of affairs. As Stalnaker puts it in the opening pages of *Inquiry*:

Rational creatures are essentially agents. Representational mental states should be understood primarily in terms of the role that they play in the characterization and explanation of action. What is essential to rational action is that the agent be confronted, or conceive of himself as confronted, with a range of alternative possible outcomes of some alternative possible actions. The agent has attitudes, pro and con, toward the different possible outcomes, and beliefs about the contribution which the alternative actions would make to determining the outcome.... This picture suggests that the primary objects of attitudes are not propositions but the alternative possible outcomes of agents' actions, or more generally, alternative possible states of the world. (1984, p. 4)

Most accounts of belief (and not merely a possible worlds account) are ones on which belief divides the space of possibilities into two. A subject who believes that p , for instance, rules out as actual all possible worlds in which it fails to be the case that p . What is distinct about a possible worlds account of belief, however, is the idea that an agent's attitudes are primarily directed toward possibilities. The content of a belief is *given by* the division it makes in the space of possibility.

This pragmatic picture of belief, on which propositional attitudes are directed toward possibilities, contrasts with what Stalnaker calls the linguistic picture, according to which representational content (of belief, desire, and other propositional attitudes) is sentence-like.² On this kind of model, belief states represent the world because they resemble, or because

² See, however, Currie 1986, p. 569 for a complaint about Stalnaker's calling his picture, in contrast to the linguistic picture, pragmatic.

they are related in a certain way to, linguistic expressions. Whereas in a possible worlds model belief states are principally directed toward possibilities, according to the linguistic picture they are principally directed toward sentences or sentence-analogues. For instance, a view motivated by the linguistic picture might hold that a belief is directed toward a proposition understood as a structured compound, with objects, properties and relations as constituents. (Typically, such a theory would hold that the belief state itself is structured, with a structure that parallels that of the proposition it has as its content.)

We need not consider in any detail the alleged theoretical advantages of possible worlds models, or the difficulties they face, since what is at issue is not whether such a model of belief is correct.³ The goal is only to assess whether either Stalnaker or Lewis offer a view on which a subject has a single, overall representation, and one that provides an adequate account of what it is for a subject to have a single, overall representation. This task we can undertake without assessing whether we should adopt a possible worlds semantics of belief.

Lewis's Modal Semantics of Belief

Lewis modifies the basic framework of Stalnaker's account, but the picture is substantially the same. He agrees with Stalnaker that belief content should not be understood by appeal to quasi-linguistic models, adverting to Stalnaker for a defense of the general possible worlds model of belief. (1986, p. 28) "A more promising plan", he says, "is to characterize the content of knowledge or belief from the outset in terms of something like the epistemically or doxastically accessible worlds." (1986, p. 28) These notions he introduces as follows:

³ See Stalnaker 1984, Chapter 4, for an overview of alleged benefits, and Robbins 2004 for a compact presentation of some of the central problems. See Soames 1987 for a more detailed attack on the possible worlds account of propositions.

The content of someone's knowledge of the world is given by his class of *epistemically accessible* worlds. These are the worlds that might, for all he knows, be his world; world W is one of them iff he knows nothing, either explicitly or implicitly, to rule out the hypothesis that W is the world where he lives. Likewise the content of someone's system of belief about the world (encompassing both belief that qualifies as knowledge and belief that fails to qualify) is given by his class of *doxastically accessible* worlds. World W is one of those iff he believes nothing, either explicitly or implicitly, to rule out the hypothesis that W is the world where he lives. (1986, p. 27, italics in original)

As in Stalnaker's proposal, belief content is "given" by classes of worlds, or classes of worlds characterize the content of belief "from the outset" (Lewis 1986, p. 28). That is, it is not that something else, such as one's having certain sentences in a language of thought in one's brain, or standing in a certain kind of relation to a structured proposition, or having accepted certain sentences in one's language, etc., originally characterizes the content of belief, after which one can speak of worlds in which one's beliefs are true and worlds in which one's beliefs are false. It is that the content of belief states just is to be understood by appeal to how the belief state divides the space of possibility.

Lewis agrees with Stalnaker that whatever else fixes or determines content, when it comes to providing a specification of what it is that a subject believes, the way to do that is by appeal to classes of worlds. He adds, however, that what has content is *a system of belief*.⁴ Lewis says:

If content is given by a class of doxastic alternatives (or by a probability distribution)⁵, what is characterized is one whole system of belief, not several beliefs

⁴ I have simplified Lewis's account, for he introduces a complication to accommodate *de se* belief, which he takes to be irreducible to belief about the world. Roughly, the idea is that knowing everything there is to know about the world doesn't entail that one knows which individual one is within the world. So, instead of saying that the content of a subject's beliefs is characterized by the class of doxastically accessible worlds, Lewis says that the content of a subject's beliefs is characterized by the class of doxastic alternatives – those individuals who might, for all that S believes, be S. For instance, my belief that *p*, is given by the class of individuals such that nothing I believe rules out that I am one of them. These individuals have various properties, of course, among which (for Lewis) are the properties of inhabiting worlds in which such and such are true. Thus, my belief that *p* is characterized by the class of individuals who, for all I believe, may well be me, and all of those individuals inhabit worlds in which *p* is true. This refinement does not have any bearing on my discussion of the issues here, however, so for the purposes of simplicity I work with the simpler variant.

⁵ Here Lewis appeals to a second refinement of the basic view, which I have not considered. No harm is done by working with the simpler view.

– the relevant notion of belief is singular, not plural. This built-in holism⁶ is one way in which the present approach contrasts with strategies in which there is a different belief for every different sentence of the language of thought that is written in the ‘belief box’. There is no sensible question whether something is one of your beliefs in its own right, or whether it is merely a consequence of some of your other beliefs. There is no sensible question whether your belief that you are hirsute is or isn’t the same belief that you are hairy; your doxastic alternatives are all hairy, in other words they are all hirsute; and that’s that. What is written in your ‘belief box’, if anything, or what word if any you might use to express yourself, is beside the point. (1986, p. 32)

There are two points in the above passage that we should be sure to distinguish, each of which might be thought to go some way toward supporting belief monism. The first is the claim that what is characterized is one whole system of belief, not several beliefs – the relevant notion of belief is singular, not plural. If this is right, then it seems that what a believer has is something like one belief that represents the world as being some particular way. This would be, it is very natural to think, one (overall) representation of the world. Hence it appears, on the face of it anyway, that Lewis is arguing that a subject has a global representation of the world and offering an account of what such a representation is and of what makes it the case that a subject has such a representation. The second point is the claim that the account differs from language-of-thought accounts by rendering unintelligible traditional thorny problems concerning the individuation of belief. This too might reasonably be seen to support belief monism, for it puts critical pressure on the notion of ‘individual belief’, without which, it seems to me, it cannot be the case that a subject has many beliefs.⁷

Let’s begin by focusing on the first point, according to which what is characterized is one whole system of belief and thus (allegedly) belief is singular rather than plural. It is

⁶ Lewis calls his view *holism*, but I think it better to call the view *monism*. ‘Belief holism’ is usually reserved for one of two distinct theses: (i) necessarily, if a subject has any beliefs, she has many beliefs; and (ii) the content of beliefs is in some sense interdependent or mutually determined. (See Chapter Three: “Davidson’s Holism”, for more discussion.) This is not what Lewis is claiming. His view is rather that when a subject believes she has what is best thought of as *one* belief. Hence ‘monism’ is a better term than ‘holism’.

⁷ Here I agree with a remark made by Leibniz in a different context: “[W]here there is no true unity, there is no true multiplicity.” Quoted in Russell 1900, p. 107.

unclear why Lewis thinks that what has content is (only) a subject's entire belief system, and thus it is unclear why, even if we characterize belief content by appeal to classes of possible worlds, we are bound to think of belief as singular rather than the plural. In *On the Plurality of Worlds* Lewis offers us no reason to believe that in characterizing belief content we need to characterize a subject's "whole system of belief", construed as one thing. And this is puzzling, for it seems we could just as well characterize sub-systems of an individual's whole system of belief, or individual constituents of an individual's whole system of belief.

By way of analogy, we could characterize the content of the *Encyclopaedia Britannica* by appeal to a class of worlds – those worlds in which everything in the *Encyclopaedia Britannica* is true. Because there are some questions on which the *Encyclopaedia Britannica* is silent, there will be more than one world in this class. It is true that *if* we characterize *all* of the *Encyclopaedia Britannica* by appeal to a class of worlds, the relevant characterization is singular, or a characterization of the content of an individual (*the EB*). But we could also characterize the content of individual volumes within the *EB*, of individual entries, or even of individual sentences. To characterize, for instance, the content of individual entries would be to assign to each entry a class of worlds – all those worlds in which what that entry says is true. Some of those worlds will be worlds in which what some other entries in the *EB* say is false, while others will be worlds in which what all entries in the *EB* say is true. It is true that if we set out to characterize the content of the *EB* by appeal to worlds, we end up with one class of worlds characterizing the content of one thing. But we are free not to do that: we could just as well characterize the content of individual entries, or of individual sentences, in the *EB* by appeal to classes of worlds, and then characterize the content of the whole by appeal to the intersection of those classes.

Much the same would seem to be true for belief. If we characterize the content of a subject's "whole system of belief", the relevant notion of belief is singular rather than plural. But that is because what we have set out to characterize is *an individual*, the system of belief. Similarly, if we characterize content by appeal to S's doxastic alternatives, where these are defined as those people who, *for all that S believes*, could be S, what we characterize is the content of one thing, S's whole system of belief. But again this is because we are characterizing content by appeal to *all* of what S believes; it should be no surprise that this means we are characterizing the content of S's entire belief system. What is missing from Lewis's account is a reason to believe that this is the way we ought to proceed. This is unfortunate, for his claim in *On the Plurality of Worlds* that what has content is principally, or even really, an entire system of belief is repeated elsewhere, and seems to be something he thinks is important. In "Reduction of Mind", for instance, he says "The contentful unit is the entire system of beliefs and desires" (1999, p. 324), and he affirms that he is a "robust realist...[a]bout whole systems of beliefs and desires, anyway, though maybe not about all the little snippets – the sentences written in the belief and desire boxes – of which these systems may or may not be composed." (1999, p. 322) For all that Lewis appears to *argue*, however, we could just as well set out to characterize the content of individual belief states within that system of belief, and then read off the content of the whole system by appeal to the intersection of the classes of worlds that characterize each individual belief. It does not seem, therefore, that we are inevitably led to think of belief as singular rather than plural by committing to characterizing content by appeal to classes of possible worlds. If belief is singular rather than plural, or singular in addition to plural, it must be for some other reason.

Although Lewis does not appear to explicitly specify why it is that what has content is only an entire system of belief, it may be possible to reconstruct some of what he has in

mind. Two general themes emerge, both in *On the Plurality of Worlds* and elsewhere. The first is an emphasis on what Lewis calls constitutive rationality. The second is a commitment to functionalism. Let us examine whether these might support the claim that what has content is principally an entire system of belief.

Lewis, like Davidson, thinks that rational or logical relations between beliefs (and between beliefs and other propositional attitudes) are in some sense constitutive of them. He says:

If the content of belief...is not tied in any uniform and straightforward way to the truth of ordinary language ascriptions of belief, and also is not tied to the subject's acceptance of inner sentences, how is it tied down at all? I would say that it is tied down mainly by belief-desire psychology. We suppose that people tend to behave in ways that serves their desires according to their beliefs. We should take this principle of instrumental rationality to be neither descriptive nor normative but *constitutive* of belief. It enters into the implicit definition of what it is for someone to have a certain system of belief. (1986, p. 36, italics in the original)

The same point is pressed in "Reduction of Mind", a 1994 paper (Lewis 1999 in Bibliography) that summarizes his views of the mental; there he emphasizes (and gives a name to) "constitutive rationality", the claim that a subject's propositional attitudes don't merely happen to be largely rational, but are necessarily so, and moreover that rational connections play a constitutive role in determining the content of doxastic and volitional states. Lewis does not spell this out in great detail, nor so far as I know has this been done in the secondary literature. I take it that the idea is very similar to Davidson's interpretive or inter-attitudinal holism. What make a belief the belief it is are the rational relations it stands in to a subject's other beliefs and desires. What Lewis then adds is that the assignment of content is an assignment of content to one state, rather than a simultaneous or coordinated assignment of content to many distinct states. He says:

I have been speaking as if the assignment of content were an assignment directly to a given subject. But I would rather say that the content belongs to some state – a brain

state, perhaps – that recurs in many subjects....The recurrent state would tend to dispose anyone who had it to behaviour fitting a certain reasonable assignment of content. Therefore we can say that the state *is* a system of belief and desire with that content, and when a subject has that state he thereby has the content that belongs to the state.⁸ (1986, p. 39)

This quotation illustrates how Lewis adds to the basic claim that rational relations are constitutive of the content of propositional attitudes the additional claim that what has content is some total state. It also illustrates the functionalism that motivates Lewis's approach, and which is emphasized by Stalnaker in describing the possible worlds account of intentional content as being pragmatic. Whereas interpretive or inter-attitudinal holism, and presumably constitutive rationality, tends to emphasize logical or rational connections between propositional attitudes, functionalism emphasizes causal connections. Hence a second idea that may be driving Lewis's commitment to characterizing the content of the whole system of beliefs is something like the following: beliefs, desires, intentions, etc., are states defined by their functional role, but a functional role, by definition, is a role that is specified by appeal to entire causal system. (For a defense of both points, see Lewis 1970 and 1972, in which he proposes that the causal system should be given by Ramsification of the platitudes of folk psychology⁹). Hence what a subject believes is determined by the entire functional state of the subject. This functional state is an internal state of a subject that disposes a subject to behave in certain ways in certain circumstances, and that "fits" some reasonable assignment of beliefs and desires.¹⁰

⁸ The motivation here is as follows: "The reason why I prefer to attach content to the state, rather than directly to the subject, is that it leaves room for exceptional cases in which, despite the constitutive role of principles of fit, the subject's behaviour somehow fails to fit his system of belief and desire. I said that the state *tends* to dispose anyone who has it to behave in a certain way; but such a tendency might be defeated." (1986, p. 39)

⁹ This involves (i) construing the content of folk psychology as a single conjunctive sentence, (ii) substituting free variables for the theoretical terms such as belief that *p*, desire that *q*, etc., and then existentially quantifying the sentence. The term comes from the surname of Frank Ramsey, who offered the model as an account of the meaning of theoretical predicates (Ramsey 1931).

¹⁰ The full story is a bit more complicated, for Lewis thinks many assignments of beliefs and desires fit the dispositional state equally well, but that certain assignments are less reasonable or appropriate. See Lewis 1986, p. 38 for the details.

I should emphasize that in sketching how constitutive rationality and a commitment to functionalism might ground Lewis's claim that what has content is principally a system of belief, I am not claiming that this is definitely what Lewis has in mind. Rather, what I am trying to do is sympathetically construct some support for the claim he makes.

That said, I confess that I do not see how either constitutive rationality or Lewis's analytic functionalism entails that what has content is only the entire belief system. It seems to entail only something like this: a propositional attitude type (a belief with content C, a desire with content C, etc.) is, necessarily, tokened only if appropriately many other propositional attitudes types are also tokened, where which other propositional attitude types are "appropriate" will depend on the particular attitude type in question, and on whether one is emphasizing rational/conceptual relations (as in constitutive rationality) or causal relations (as in functionalism). Both constitutive rationalism and Lewis's analytic functionalism, to be sure, account for why beliefs and desires cannot be had on their own (why a belief or desire cannot be had on its own, independent of a whole system of beliefs and desires). But this does not seem to require that we say that what has content is either principally or only an entire system of belief. This is perhaps easiest to see in the case of functionalism. If we follow Lewis and imagine a theory as a Ramsified sentence with variables in place of the theoretical terms, it may be that the meaning of the theoretical terms is given by the entire theory in which they feature. But it still seems that on this account the theoretical identities hold between the causal roles so specified and whichever entities fill those roles. The causal roles and hence the entities may be specified all at once, but this doesn't entail that what gets specified is one thing. Remember, too, that that analytic functionalism is an account of theoretical terms generally, not merely of mental predicates. So, for instance, it plausibly provides an account of the meaning of theoretical terms and identities in physics. Yet it is

hard to see that a proponent of such an account would be committed to saying that it makes no sense to speak of force and of mass and of acceleration, merely on the grounds that the terms 'force' and 'mass' and 'acceleration' are given by appeal to an entire theory.

Individual Beliefs

The second point Lewis makes in the passage quoted above is that what he calls the holism of his account distinguishes it an important way from accounts which take the having of a belief to be the having of a sentence in a language of thought within one's 'belief box'. He says:

There is no sensible question whether something is one of your beliefs in its own right, or whether it is merely a consequence of some of your other beliefs. There is no sensible question whether your belief that you are hirsute is or isn't the same belief that you are hairy; your doxastic alternatives are all hairy, in other words they are all hirsute; and that's that. What is written in your 'belief box', if anything, or what word if any you might use to express yourself, is beside the point. (1986, 32)

This is worth considering briefly, since what Lewis challenges is the intelligibility or correctness of the notion of individual beliefs, which might be seen to go some distance towards supporting belief monism. In my view, Lewis is running together two distinct claims. One is that there is no sensible question whether the belief that *p* is distinct from the belief that *q* where '*p*' and '*q*' are synonymous terms. The second is that there is no sensible question whether someone who believes (for instance) *that p* and believes *if p then q* also has the belief *that q* "in its own right".

Consider Lewis's claim that there is no sensible question whether your belief that you are hirsute is or isn't the same as your belief that you are hairy. One might think Lewis holds that this is not a sensible question because there is no difference between being hairy and being hirsute – as he says, "your doxastic alternatives are all hairy, in other words they are all hirsute, and that's that." But if this is his reasoning, then it seems the question of

whether your belief that you are hirsute is or isn't the same as your belief that you are hairy *is* a sensible question – the two beliefs are really one, for there is no difference between one's doxastic alternatives all being hairy and one's doxastic alternatives all being hirsute.

Alternatively, one might think that Lewis holds that there is no sensible question whether one's belief that one is hirsute is or isn't the same as one's belief that one is hairy because of his monism, according to which there are not really any such things as individual beliefs – there is only the whole belief system (as he says, the relevant notion of belief is singular not plural). This promises to make sense of how the question under discussion can be without sense, for if there are no such things as the individual belief that one is hirsute or the individual belief that one is hairy, then it makes no sense to ask whether the two are the same or different. But if this is what Lewis has in mind, he chose his example poorly, for by talking about 'hirsute' and 'hairy' it seems as if what he thinks renders the question nonsensical is the fact that the terms are extensionally equivalent. If the question is nonsensical because there are not really any such things as individual beliefs, Lewis could just as well have said there is no sensible question whether one's belief that one is hairy is or isn't the same as one's belief that Baffin Island is the fifth largest island in the world. That question isn't sensible, on this interpretation of Lewis's reasoning, because there are no such beliefs – there are only total belief systems the content of which is characterized by appeal to classes of worlds. Lewis's point does not seem to be that the question regarding the identity or non-identity of one's hairy/hirsute beliefs is nonsensical merely because there are, on his account, no such things as individual beliefs. What he seems to have in mind is that the questions are nonsensical because there is no difference between one's doxastic alternatives being hairy and one's doxastic alternatives being hirsute. But the conclusion just doesn't follow, for the reason noted above – it would seem that the proper thing to say in this case is

that the beliefs in question are identical, for they rule out the same worlds. Moreover, it is not clear how this has anything to do with the monism Lewis takes his view to involve. If belief content is given by classes of worlds, then there is no difference between believing one is hairy and believing one is hirsute regardless of whether Lewis is right that what has content is fundamentally a belief system rather than individual beliefs.

The second kind of question Lewis says his monism renders without sense are those that concern whether something is one of your beliefs “in its own right” or merely entailed by your other beliefs. Lewis does not make his reasons for holding this view perspicacious. One option, as before, is to take the point to be that on his view there are no such things as individual beliefs and so, by virtue of that, no sensible question whether something is a belief in its own right (and therefore, *a fortiori*, no sensible question whether something is a belief in its own right or merely a consequence of other beliefs). But this is unsatisfying for the same reasons given above. If this were what Lewis had in mind, there would be no point in specifying that the question rendered unintelligible by his view of belief content concerns beliefs that are entailed by one’s other beliefs. The idea seems to be, rather, that the question is nonsensical because all worlds in which p and *if p then q* are true are worlds in which q is true, and thus whenever one believes p and *if p then q* the class of worlds which characterizes the content of one’s belief system is the class of worlds which would characterize the content of one’s belief system were one to believe, in addition to what one already believes, q . (The idea here being that there is no difference between believing $\{p, \text{if } p \text{ then } q\}$ and believing $\{p, \text{if } p \text{ then } q, q\}$ – the same class of worlds characterizes each believing and thus, on Lewis’s view, the content of each believing is identical.)

It is unclear, however, how this is supposed to show that it is nonsensical to ask whether something is one of your beliefs “in its own right” or merely entailed by some of

your other beliefs. It is true that if content is characterized by classes of worlds, then it isn't possible to believe $\{p, \text{if } p \text{ then } q\}$ without believing $\{p, \text{if } p \text{ then } q, q\}$. But this does not show¹¹ that there is no such thing as believing a proposition in its own right – for one could just as well say that it is impossible to believe propositions $p_1 \dots p_n$ while failing to have the belief in its own right that q , where q is a proposition entailed by $p_1 \dots p_n$. What we need, again, is a reason to think that in characterizing content by classes of worlds we are required to characterize only the content of an entire system of belief.¹²

I have detoured briefly through these issues because what is at stake is the intelligibility of the notion of an individual belief. I take it that Lewis, in the passage quoted, is skeptical about this idea, and one might think that if the skepticism is well-founded it goes some way toward supporting his belief monism. Unfortunately, it is not clear how the remarks he makes are supposed to show that there is anything remiss with the notion of an individual belief.

Compartments Again

We do not find in Lewis compelling reason to accept that content attaches principally, let alone only, to entire belief systems, nor, as I see it, good arguments against the view that it makes sense to speak of individual beliefs. There are two further points to be made

¹¹One might insist that if beliefs are real entities, then for every proposition p either the subject has the belief that p or fails to have the belief that p , and, further, that in the case I am imagining there is no fact of the matter whether the person has the belief that p . In response I would ask why we are compelled to think that in the case imagined there is no fact of the matter whether the person has the belief that p . If anything, it seems that on the model of content we are supposing is true there is a fact of the matter about whether the person has the belief that p – she rules out as actual every world in which $\sim p$, which is just what it is to believe that p .

¹²I am not sure how to reconcile Lewis's claim that an advantage of a possible worlds semantics of belief content is that there is no genuine question concerning whether one believes a proposition that is entailed by one's other beliefs, with his claim, discussed below, that his view does not entail that subjects are logically omniscient, believing all the logical consequences of their beliefs. He seems to want to say both that there is no difference between believing a set of propositions and believing that set of propositions plus the propositions those propositions entail, *and that* subjects do not always believe the consequences of their beliefs (such as in those cases where the entailing beliefs are held in different mental 'compartments'). This seems straightforwardly inconsistent.

concerning whether Lewis's theory of belief content is on one which, in general, a subject has a single, overall representation. Lewis advertises his theory as a monistic account (although he does not use that term), according to which belief is singular rather than plural. As I have argued, it is unclear why Lewis thinks that what needs to be characterized is the content of an entire system of belief, rather than of subsystems of belief, and thus it is unclear why a commitment to characterizing content by appeal to classes of worlds leads one to belief monism. I now want to note that even if these arguments fail and we accept, with Lewis, that content attaches in the first instance to belief systems, rather than to individual beliefs within such systems, his view is still not properly thought of as monistic. This is because Lewis retains plurality but locates it in *systems of belief* rather than in beliefs. That is, although he says that on his view the relevant notion of belief is singular rather than plural, on the grounds that what is characterized as having content is an entire system of belief, he holds that a subject has numerous belief systems at the same time. He does this in order to handle problems that otherwise arise for his view that content is characterized by classes of worlds.¹³

I'll consider just one of those problems, to give a feel for his proposal. The first problem is that of what Lewis calls "double-thinking". He says:

[S]omeone may have a multiplicity of belief systems. To a greater or lesser extent, we are all doublethinkers: we are disposed to think differently depending on what question is put, what choice comes before us, what topics we have been attending to. Belief is compartmentalized and fragmented. Sometimes a doublethinking believer acts in a way that best fits one belief system, sometimes in a way that best fits another. And it should not be said just that his belief system changes rapidly; because, throughout, he remains simultaneously disposed toward both systems. (1986, pp. 30-31)

¹³ In this regard Lewis is similar to Davidson (see Chapter Three: Davidson's Holism). It should be noted that Lewis is here following Stalnaker, who invokes talk of compartmentalization and belief fragmentation in Chapter V of *Inquiry*, titled "The Problem of Deduction". I here set out Lewis's account of fragmentation by quoting from his 1986, but see also Lewis 1982.

He illustrates the claim with an example of a person who is simultaneously disposed to both hypochondria and good cheer:

You have one belief system, the hypochondriac one, under which all your doxastic alternatives are in the early, invisible stages of a dread disease. You have another belief system, the cheery one, under which all your alternatives are healthy. (1986, p. 31)

Lewis takes pains to emphasize this is not a case of a person who is unsure whether he is healthy or ill: such a person, as Lewis notes, assigns some credence to both the proposition that he is healthy and the proposition that he is ill. Moreover, if you ask such a person whether he is ill or healthy, he'll likely respond in a way that makes clear that he is uncertain. Rather, the case in question involves a person who, as Lewis somewhat obscurely puts it, "half-heartedly certain" he is healthy and "half-heartedly certain" that he is ill. If you were such a person and were asked to bet on whether you are healthy or ill, "which way you go depends on how exactly the question is put to you, and on how you're feeling at the time." (31)

Such a person, on Lewis's view, does not have *a* belief system – he has at least two, at the same time. Presumably, each case of doublethinking would multiply the number of systems of beliefs had by a person by two, such that if the hypochondric/good cheer person also is a doublethinker in regard to whether the Mets will win the pennant, he has four belief systems:

BS1: [other beliefs], I'm sick, Mets will win.
 BS2: [other beliefs], I'm sick, Mets will lose.
 BS3: [other beliefs], I'm healthy, Mets will win.
 BS4: [other beliefs], I'm healthy, Mets will lose.

Of course, in some case it may be that, for whatever reason, a double doublethinking person has some of these beliefs systems but not others. It may be, for instance, that his being a doublethinker in regard to the Mets' fortunes is connected in some way with his being a

doublethinker in regard to his own fortunes, such that he has no inclination to believe both that he's healthy and the Mets lose. But even admitting such cases, it is clear that on Lewis's account we are all doublethinkers on numerous issues and thus have a great many belief systems at once.

Lewis's analysis of these kinds of cases is unsatisfying. He does not make clear why it is better to think that what is happening is that there is one believer with two different belief systems than that there is one believer who vacillates between believing that he is sick and believing that he is ill. He says that we are disposed to think differently "depending on what question is put, what choice comes before us, what topics we have been attending to", but this does not entail, at least not obviously, that we have two or more belief systems at once. It could be that what question is put, what choice comes before us, and what topics we have been attending to affect what we believe, or alter our belief reports (where what we sincerely report is not always what we believe). Perhaps the person disposed towards both hypochondria and good cheer has two different belief systems. But an equally plausible explanation is that he sometimes believes he is ill and sometimes believes he is healthy, depending on what question one asks, what he has been thinking about, how his day has been going, and so on. Or that his belief state concerning illness is one of uncertainty, but how the question is put, what he has been thinking about, etc., induces him to report inaccurately on what he believes. It might even be – for all that Lewis has established – that he believes both that he is healthy and that he is not healthy, and that how the question is put, how his day is going, etc., makes one belief more accessible to reporting than the other. Lewis's account might be the best of the four – but he has not given us good reason to believe that it is.

Lewis's account also raises some very puzzling questions, questions one might think it would be better not to have to answer. On his account, a person somehow doubles the number of belief systems she has simply by being a 'double-thinker' in regard to some issue. Suppose I am not a double-thinker at all, and that my believing is not 'fragmented' in any other way. I have one belief system. By double-thinking on the issue of my health, I double the number of belief systems I have – from one to two. If I double-think in regard to something else, I now have four belief systems instead of two. My worry about this is that I have very little sense of what it means. What is it that doubles? What are there now twice as many of? By what process does this doubling occur? Lewis's proposal has the air of the magical, and hence of the unreal.

Whether or not Lewis is right to claim that a subject has a multitude of belief systems simultaneously,¹⁴ however, it seems clear that this claim rules out a reading of Lewis on which a subject's believing is, at a time, genuinely one. He describes his account of belief as one on which "belief is singular rather than plural", and for this reason his theory seems to commit to the monistic claim that a subject's believing is one. His talk of multiple belief systems, however, makes it evident that he retains plurality but locates it in systems of beliefs rather than in beliefs themselves. For this reason, his account is not properly thought of as

¹⁴ Lewis has other reasons for thinking that belief is compartmentalized or fragmented. As one example, however:

Suppose that you believe P, also you believe Q, and P and Q jointly imply R in the sense that every world that is both a P world and a Q world is also an R world; nevertheless, we may suppose that you fail to believe R. We may even suppose that *none* of your doxastically accessible worlds is an R world. How can this be? – The answer is that you may be thinking double, with P and Q in different compartments. You believe that P by believing it in one system; that one gives you doxastically accessible worlds where P holds but Q and R do not. You believe that Q by believing it in the other system; that one gives you doxastically accessible worlds where Q hold but P and R do not. Thus you believe P and you believe Q, though in both cases half-heartedly; but you whole-heartedly disbelieve the conjunction of P and Q, and you wholeheartedly disbelieve R. You fail to believe the consequence of your two premises taken together so long as you fail to take them together. (1986, p. 35)

There is much about this analysis of inference (or inferential failure) that deserves question. What does it mean to say that two beliefs are in different compartments, or to say that one believes a proposition in one system but not in another, or to say one believes two propositions but fails to take them together? There seems to be little willingness to spell these notions out.

monistic, and for the same reason it is hard to see how it is an account on which a subject has *a* global representation of the world.

Stalnaker Again

As I have said, Stalnaker does not explicitly claim that his account of belief is one on which a subject has one overall belief, or one overall representational state. And he appears to be generally happy to talk of beliefs in the plural, and to apply the possible worlds model in such a way that he characterizes the content of what we usually call individual beliefs.¹⁵

However, in a number of places Stalnaker switches to talking of a *belief state*, which seems to be something that has as its content an overall representation of the world. But like Lewis and Davidson, Stalnaker adopts a compartmentalization thesis, which makes it hard to see him as arguing that a subject has a global representation of the world. In fact, he appears to pointedly deny that a subject has such a representation, and the denial amounts to his offering a condition on a subject's having a global representation of the world.

In a number of places Stalnaker seems to take the notion of a belief state to be the fundamental doxastic element, where a belief state is something like a conception of the world that in some sense subsumes or contains what we call individual beliefs. For instance, he writes:

¹⁵ For instance, in "Assertion", Stalnaker writes "A proposition – the content of an assertion or belief – is a representation of the world as being a certain way. But for any given representation of the world as being a certain way, there will be a set of all the possible states of the world which accord with the representation – which *are* that way. So any proposition determines a set of possible worlds. And, for any given set of possible worlds, to locate the actual world in that set is to represent the world as being a certain way. So every set of possible worlds determines a proposition. Furthermore, any two assertions or beliefs will represent the world as being the SAME way if and only if they are true in all the same possible worlds. If we assume, as seems reasonable, that representations which represent the world as being the same way have the same content (express the same proposition), then we can conclude that there is a one-one correspondence between sets of possible worlds and propositions. Given this correspondence, it seems reasonable to use sets of possible worlds, or (equivalently) functions from possible worlds into truth values, to play the role of propositions in our theory. The analysis defines propositions in terms of their essential function – to represent the world." (2006, p. 197, italics and all-capitals are Stalnaker's)

A belief state can be represented as a set of possible worlds. Individual beliefs are properties of such a belief state: to believe that P is for the proposition that P to be true in all the possible worlds in the belief state. If one conceives of beliefs in this way, they look like something negative: to believe that P is simply to be in a belief state which lacks any possible world in which P is false. To be more opinionated than one ought to be is to have too many beliefs, but having too many beliefs is the same as not recognizing enough possibilities, and according to the possible worlds conception of a state of belief, it is the latter characterization that is the more fundamental. This makes it easier to see how a person can have beliefs by default, or through a lack of imagination. And there is no difficulty, on this conception, in the idea of a finite believer with a large or even an infinite number of separate beliefs.... Just as a finite perceiver may see a space which in fact consists of an infinite number of points, so a finite believer may represent a space of possible worlds which in fact consists of an infinite number of possible worlds. (1984, pp. 68-69)

The idea presented in this passage is that the belief state, understood as a representation of ways the world could be, is itself the fundamental doxastic element. Individual beliefs are said to be “negative properties” of such a state. The suggestion is that a belief state is basic and individual beliefs are in some way derivative, for the picture is not one on which a subject has various individual beliefs that compose a belief state but rather one on which a subject represents a space of possible worlds as ways the world might be, and individual beliefs can be ‘read off’ from that. We see here, therefore, two distinct ideas: that a belief state is, as Stalnaker puts it, a “conception of the way the world is” (1984, p. 84), and that the belief state is in some sense fundamental relative to the individual beliefs it has as negative properties.

Stalnaker is no more forthcoming than Lewis when he says we should think that subjects have belief states that have individual beliefs as negative properties, rather than just that they have individual beliefs the content of each of which is given by a class of worlds. Instead, he seems to introduce the possible worlds model in schematic terms, saying that possible worlds can be used to model the content of a proposition and that a belief has as its content such a proposition. Then, suddenly, he is talking about a belief state having

individual beliefs as negative properties, with the suggestion that the belief state is the fundamental doxastic element and something like a view of the way the world might be.

This might seem close to the claim that a subject has a global representation of the world, but for the fact that Stalnaker, like Davidson and Lewis, endorses a compartmentalization thesis. In so doing, he explicitly denies that a subject has a global representation of the world, and the denial amounts to the offering of a condition a subject must meet if she is to have such a representation. The denial, and the condition, arise from the fact that on a possible worlds model of intentional content, belief is closed under entailment. In a nutshell, the idea is that although any given belief state is closed under entailment (as a defining condition of a belief state), all of a subject's beliefs are not closed under entailment because they can be, and almost always are, kept in different compartments. So one might believe that p and that if p then q , but fail to believe that q ; the explanation is that the belief that p and the belief that if p then q are in different compartments, or are 'negative properties' of different belief states. And like Lewis, Stalnaker thinks that a subject at least typically has many different belief states at once.¹⁶ He presses this point in a passage that appears to deny that subjects (or, at least, all actual subjects and all subjects relevantly like them) have a global representation of the world:

A state of belief is most perspicuously represented, not by a set of sentences or propositions believed, but by the set of possibilities recognized as ways the world might be....The beliefs of a perfectly rational intelligence could be represented by a single belief state of this kind – one coherent conception of the way the world is represented by one set of alternative possibilities. But the beliefs of mere mortals will require a more complicated representation. Mortals may be in many belief states at once, represented by separate spaces of possibilities. (1984, p. 98)

Stalnaker should be interpreted here as denying that a subject has a global representation of the world – not merely that a subject has a *coherent* global representation of the world. This is

¹⁶ One might wonder, as Gregory Currie did in a review of *Inquiry* in the *Philosophical Quarterly*, "whether Stalnaker is going to end up with belief states as finely individuated as beliefs themselves". (1986, p. 570)

because on the possible worlds model of content, there can be no incoherent representations. (Stalnaker elsewhere puts the point even more explicitly, as in the summary he provides of his view in Guttenplan's *Companion to the Philosophy of Mind*: “it must be recognized that beliefs may be fragmented: one may be in separate belief states that are not integrated into a single conception of the world” (Stalnaker 2001, p. 566).) On the picture Stalnaker is offering, therefore, it would appear that a subject has a global representation of the world only if her believing is closed under entailment. That is, to have a global representation of the world a subject would have to meet this condition:

If S believes, at t, $p_1 \dots p_n$, and p_i is a proposition entailed by any combination of $p_1 \dots p_n$, then S believes, at t, p_i .

If S fails to meet this condition, her beliefs would be compartmentalized or she would have more than one belief state, and thus no single conception of the way the world is.

This strikes me as an unreasonably high standard for a subject to have to meet in order to have a global representation of the world. Moreover, on Stalnaker's proposal, one could say of some subject S with belief states A and B something like “according to belief state A, the world is such that C”, and “according to belief state B, the world is such that C2”, but one couldn't make any sense at all of the idea that there is some way that she, the subject S, takes the world to be. Or, at least, that way wouldn't be closed under entailment, which the possible worlds model of content cannot very well make sense of (it has no resources to specify such a content).

The Map-Like Theory

Let me close with a brief look at Lewis's discussion of the map-like theory of belief, for this discussion might be seen both to independently motivate a kind of belief monism and to support the semantic argument by giving us a reason to think that what needs to be

characterized is the content of an entire belief system. I will first sketch the map-like theory and explain why one might think it offers independent reason to favor belief monism, and then explain why I think it does not. I'll close this section by discussing whether, as one might reasonably think, it could be used to support belief monism by combining with the semantic account just sketched.

In "Reduction of Mind", a 1994 paper in which Lewis summarizes his views of the mental, Lewis sympathetically presents a map-like theory of belief as an alternative to a language-of-thought hypothesis. Although he is non-committal either way, he seems somewhat more inclined toward a map-like theory. He presents it as follows:

A serious issue...concerns the relation between the whole and the parts of a representation. Suppose I have a piece of paper according to which, *inter alia*, Collingwood is east of Fitzroy. Can I tear the paper up so that I get one snippet that has exactly the content that Collingwood is east of Fitzroy, nothing more and nothing less? If the paper is covered in writing, maybe I can; for maybe 'Collingwood is east of Fitzroy' is one of the sentences written there. But if the paper is a map, any snippet according to which Collingwood is east of Fitzroy will be a snippet according to which more is true besides. For instance, I see no way to lose the information that they are adjacent, and that a street runs along the border. And I see no way to lose all information about their size and shape....

Mental representation is language-like to the extent that parts of the content are the content of parts of the representation. If our beliefs are 'a map...by which we steer', as Ramsey said, then they are to that extent not language-like. And to that extent, also, it is misleading to speak in the plural of beliefs. What is one belief? No snippet of a map is big enough that, determinately, something is true according to it, and also small enough that, determinately, nothing is true according to any smaller part of it. If mental representation is map-like (let alone if it is hologram-like) then 'beliefs' is a bogus plural. You have beliefs the way you have the blues, or the mumps, or the shivers. (1999, p. 311, *italic in original*)

Lewis is making two points about map-like representations here, although it is not clear that he is aware of the difference. One point is that maps are such that they cannot be torn up in a way that extracts exactly some single piece of information such as that Collingwood is east of Fitzroy. Inevitably, any piece of that map that represents that will represent more than that (that Collingwood and Fitzroy are adjacent, that they have a certain size and shape, etc.).

A second point is that any contentful piece of a map (any snippet big enough to be such that something is true according to it) will have a proper part that is also contentful (the piece will not be small enough such that nothing is true according to any smaller part of it). I am not sure whether Lewis is right that maps have both these features; the second claim seems especially doubtful, for it can be read to entail by regress that if x is a contentful map and y is a point-sized snippet of x , y is contentful.¹⁷ But set that aside: we can simply treat the claims made as defining what he means by ‘map-like’.¹⁸

It should be clear why this would seem to support belief monism. The idea is that if belief is map-like, then while there is a whole representation, there is no exhaustive decomposition of this whole representation into parts. (The point is not merely that there is no *unique* decomposition into representational parts; it is not that there is more than one way of decomposing the representation into parts, but that there is no way of (fully) decomposing the representation into parts, where fully here means decomposing it into smallest contentful units.) But does it? My major complaint is that Lewis appears to be talking about the *vehicle* of representation, rather than about the *content* of representation. Stalnaker makes much the same objection in his paper “Lewis on Intentionality”:

Lewis suggests that if the way information is represented is holistic – map-like, or hologram-like, rather than sentence-like – then plural propositional attitude terms such as ‘beliefs’ are ‘bogus plurals’: ‘You have beliefs the way you have the blues, or the mumps, or the shivers’ (Lewis, 1999, p. 311). But this is not right, since whatever the nature of the vehicle or vehicles of mental representation, the plural noun ‘beliefs’ does not refer to that vehicle, or to those vehicles. What it refers to is the *contents* of a representation – to the propositions that are believed. And even if there

¹⁷ Lewis does not make clear whether he is talking about ‘conceptual snippets’ or actual snippets. That is, it may be that every point or part of a map that is not yet cut up represents something, but it is very hard to imagine that every point or part represents something as a snipped snippet. The drift of the remarks Lewis is making suggests he is talking about detached proper parts of the map, and this leads to the dubious entailment.

¹⁸ In their discussion of the virtues of map-like theories vis-à-vis LOT theories, Braddon-Mitchell and Jackson characterize the map-like theory as committed only to the first of these two points, saying “maps give some (putative) information by giving a lot of information.” (1996, p. 170) It is worth noting that if semantic holism is true of whatever language features in the language of thought (assuming there is one), then this distinction between map-like theories and LOT theories collapses.

is a single map-like internal structure in a believer that make[s] it the case that she has the beliefs that she has, that structure will determine a plurality of propositions that are believed....[W]hatever the character of what is going on in a person's head when she has beliefs, these goings on should not be confused with what the person believes. (2004, p. 208, italics in original)¹⁹

The important point is that what is at issue when we are interested in whether a subject has a global representation is whether the content of a subject's beliefs is one representation, not whether the vehicle by which she believes what she believes is a metaphysical unity (i.e., one entity). Of course, it may be that the point Stalnaker is making is the point that allows us to speak of a subject having *many* beliefs, whereas the point Lewis is making, or suggesting, regarding belief being 'map-like', is what makes it the case that a subject has *one*, overall representation, understood as a claim about content. (I set aside the question of how to square the map-like thesis with compartmentalization.) I do not find this an attractive option, however, for several reasons. The fact that a vehicle of representation is map-like does not entail that instantiations of that type of vehicle at a time compose one overall representation. Maps, for instance, are map-like. But all the maps in my house – I would probably find about 20 if I looked in every drawer and box – do not compose any overall representation, or so it is natural to assume. So at the very least what we would need is reason to think not merely that belief is map-like, but that each subject has but one doxastic map. Moreover, suppose I had a map of Ontario and a map of Quebec, both at 1:50,000 scale. If I then glued the maps together (with Ontario on the left), I would have a map of

¹⁹ In *Inquiry*, Stalnaker had made clear that he took the possible worlds model of belief content to be neutral on the nature of the representational vehicle:

In attributing beliefs and desires, we are attributing certain kinds of internal causal properties which have a structure that tends to reflect the world in ways that make it appropriate to call them representations. These representations could conceivably take the form of sentences of a language of thought written in the belief center of the brain, but they also could take the form of pictures, maps, charts or graphs, or (most plausibly) a diversity of redundant forms, none of which are very like any of the forms which our public representations take.

The pragmatic picture and the possible worlds definition of proposition does not then deny that beliefs are internally represented. But it remains neutral on the form that those representations must take....The aim of the definition is to give an account of the structure of what is represented while leaving open questions about the means by which this is accomplished. (1984, p. 22)

Ontario/Quebec; one representational vehicle that has, it seems right to think, one representation as its content. But the mere fact that the two maps are glued together (i.e., into one entity) is not that in virtue of which the content it has as its content is one representation. Rather, the unity belongs to the *type* of content represented (i.e., anything that represents *that content* has one representation as its content). For this reason, to understand the unity of belief we need to look not to the unity of the vehicle, but to the unity of content itself.

Conclusion

Lewis claims that belief is singular rather than plural, but he retains plurality and merely locates it elsewhere. Stalnaker's account of belief is similar, but without the monistic billing – it is an account on which a belief state is something like one representation, but it fragments a subject's believing into some unknown number of belief states. Neither account, I believe, really explains why we should think of content as something had principally by belief systems (as in Lewis) or by belief states (as in Stalnaker), rather than by individual beliefs. Stalnaker's account makes a positive contribution to the general issue of understanding global representation by offering a condition a subject has to meet if she is to have such a representation (her believing has to be closed under entailment), but as I have said it seems to me he sets the bar implausibly high.

Before I turn to developing a positive proposal, I want to draw a lesson from the discussion in this chapter that Stalnaker's complaint about Lewis's saying belief is a 'bogus plural' helps make clear. What is at issue is not whether the vehicle of representation is unified, but rather whether the content is. And it would seem that whenever a vehicle has as its, or vehicles have as their, content one representation, this is by virtue of having as content

a type, where the type is that of a unified representation with such and such content. The task, then, is to understand the nature of that type.

UNITY AND CONJUNCTION

[O]ne can never fully disentangle questions about the nature of representation from questions about the nature of what is represented.

– Robert Stalnaker, to the photographer Steve Pyke, 2003

The discussion in the previous chapters has taught us what a global representation of the world is not. It is not the same thing as a coherent belief system, nor something that can be understood by appeal to a certain kind of cognitive architecture, nor something that the massive modularity thesis in empirical psychology entails a subject cannot have. It is not something the existence of which follows from belief holism, semantic holism, or interpretive holism, nor does a subject's having such a representation follow from either a possible worlds semantics or a commitment to a map-like theory of belief. What then is a global representation of the world, and under what conditions, or in virtue of what, does a subject have such a representation?

This chapter presents and considers a two-part suggestion concerning how we may be able to understand global representation and the unity of belief. First, I attempt to formalize what I take to be the core, intuitive idea of an overall global representation (where, as set out in Chapter 1, a global representation of the world is an overall global representation whenever the overall global representation has as its representational target the world).¹ I then show how what is perhaps the most natural way to understand global representation – an appeal to conjunctive belief – is an instance of this schema. I then raise a

¹ As I said in Chapter 1, I think a subject's overall global representation is, necessary, a global representation of the world, but since this could be disputed it is important that we distinguish them conceptually.

couple of concerns I have with this two part approach, where the drift of the discussion is towards skepticism concerning an appeal to conjunctive belief. As a whole, then, the chapter explores an important model for understanding the unity of belief, trying to draw out some of its virtues, but concludes by raising some concerns with the adequacy of this model.

What is an overall global representation?

The basic, intuitive idea behind that of an overall global representation is that it is an overall picture, where what is true according to this picture is all and only what is true according to one's individual beliefs. There is plurality – a subject believes many distinct things to be true, or has many distinct beliefs. And there is unity – there is some one way the subject believes things to be, or she subject has a single, overall view. A natural thought is that the core idea is that a subject's overall global representation is a representation that in some sense contains, or has exhaustively as parts, the content of each of the subject's individual beliefs. But this needs to be made more precise: if we are trying to understand what it is for a subject's many beliefs to compose one overall representation, it will not do to use mereological terms, or composition/constituent talk, in the analysis, unless those terms are themselves defined.

As a first pass, one might try to give more precise content to this natural thought as follows:

A subject S has an overall global representation iff there is a proposition p such that S believes p and $\forall x$ (x is a proposition that S believes $\rightarrow p$ entails x)

This formulation attempts to cash out the relevant notion of 'containment' by appeal to the more precisely understood relation of entailment; the intuitive idea that the subject has a representation of the world that contains, or has as a part, everything that she believes is

analyzed as the claim that she believes some proposition that entails every proposition that she believes. This, one might hope, captures both the unity and the plurality, for on this picture the subject believes a single proposition that entails each of the many propositions that she believes.

This formulation is not quite enough, however. Presumably, on this model it is not merely that the subject has an overall global representation iff there is a proposition p as described, but that, in addition, p gives the content of the subject's overall global representation. The problem, therefore, is that one might reasonably think there is a possible situation in which a subject believes some proposition p that entails every proposition she believes, but that also entails some proposition she does not believe. Suppose, for instance, p entails both q and r , and the subject believes r but not q . In this case the intuitive thing to say is that the subject's representation of the world 'contains' r but not q , on the grounds that the subject believes one but not the other. But the formulation just sketched does not appear able to make sense of this. In regard to the proposition p that is here supposed to be the content of the subject's global representation of the world, proposition q and r are on equal footing (i.e., p entails both q and r). How could it be then that the subject's global representation of the world 'contains' one but not the other?

The point being made here could be easily misunderstood. I am not supposing that if p is the proposition that is the content of the subject's overall global representation, then the subject should believe every proposition entailed by p , and then pointing out that the subject may fail to believe some proposition r entailed by p . The point is rather that on the model sketched, the content of the subject's overall representation is p . Intuitively, since the subject believes r but not q , the subject's overall global representation should contain, in however that should be understood, r but not q . Yet r and q stand in the same relation to p (that of

being entailed by it). How could it be, then, that they differ in terms of whether they are contained by it? Here is another way of putting the point: If the model sketched were correct, then if the subject as imagined, who does not believe q , were to come to believe q , the content of her overall global representation would not change – for on this model p is as suited to be the content of the subject's overall global representation of the world after the subject believes q as before.

Another problem with this account is that there might be cases in which a subject believes a proposition p that entails every proposition that she believes, but in which the entailment between p and some proposition that she believes is such that it does not seem right to say that p contains it, at least where that vague term captures something important about the notion of a global representation. Suppose, for instance, that it is metaphysically necessary that every object that has a color has a shape, and suppose I have (per impossible) just two ‘individual beliefs’, that everything that exists has a color, and that everything that exists has a shape. If, as we are assuming, it is metaphysically necessary that every object that has a color has a shape, then the content of the first belief, that everything that exists has a color, entails the content of the second belief, that everything that exists has a shape. Thus, if the general formulation above were correct, if I believed just that everything that exists has a color and that everything that exists has a shape, then I would have an overall global representation *just by virtue of believing that everything that exists has a color*, for that proposition is both (i) a proposition I believe, and (ii) such that it entails every proposition I believe. This example is a little strained, but it suffices, I think, to illustrate an important point: Entailment may be necessary for containment, but it isn’t sufficient, for if a proposition believed is to be contained within an overall global representation, it has to in some sense be there to be seen

within the representation. Another way of putting the point is this:² If I know what your overall global representation is, I should know what your beliefs are. But knowing that you know that everything has a color doesn't tell me anything about whether you believe that everything has a shape, or even whether you have the concept 'shape'.

As another way of dramatizing the problem with this first model, consider a painting of Saul Kripke giving a lecture. If we imagine the content of the painting as a proposition, then, on the assumption of a familiar and plausible metaphysical doctrine defended by Kripke, that proposition entails that Kripke's parents exist (or existed) at whatever world the lecturing occurs. But that is surely not a proposition 'contained' by the painting, unless of course it includes a depiction of his parents in the audience, etc. The problem with the first schematic formulation is, therefore, that the reach of entailment is both too far and too indiscriminate.

These worries suggest a number of possible refinements. One might appeal to the notion of the *a priori*:

A subject S has an overall global representation iff there is a proposition p such that S believes p and $\forall x$ (x is a proposition that S believes $\rightarrow x$ can be derived a priori from p)

This retains entailment, but adds a further condition designed to capture the idea that the overall proposition that is the content of the subject's global representation of the world does not merely entail every proposition she believes, but also contains it in an appropriate sense. (I am here supposing that there is something like an absolute, idealized notion of a priori derivability, and further, of course, that the set of propositions that can be derived a priori from a proposition p is a proper subset of the set of propositions entailed by p .) This refinement is not satisfactory, however. It does not appear to respond at all to the worry

²Thanks to Jaegwon Kim for this formulation.

raised above, for under this refinement it is still the case that whether or not a given proposition is 'contained' by an overall global representation is not determined strictly by the content of that representation, but only by that content in concert with something else (with whether the subject believes the proposition in question). This is so because, where p is the content of the global representation, it seems there could be cases in which both q and r are derivable *a priori* from p , but only r and not q is believed by the subject; intuitively, r but not q is 'contained' by the subject's global representation, but they are on equal footing in regard to p , the propositional content of that representation. Once again, (i) whether a proposition is 'contained' by an overall global representation would not be determined by the proposition that is the content of that representation, and (ii) two subjects could have an overall global representation with the same proposition as its content even though one believed a proposition the other did not. This cannot be right.

The shift from entailment to the *a priori* is well motivated in one sense, however, for it is a shift from understanding the relevant sense of containment by appeal to the extensional, to a relation that holds between the entities that are represented, to understanding containment by appeal to the intensional, to a relation that hold between entities as represented. (I am here supposing, without discussion, that even the idealized *a priori* derivability is essentially intensional.) Consider the entity E1, which is the bicycle in my front yard, and the entities E2...En, which are such things as its wheels, wheel bearings, frame, bottom bracket, and so on. E1 contains or is composed of E2...En. But when it comes to representation, it seems obvious that one can represent E1 without representing any of E2....En, or at least without representing all of them. The point is that when specifying what is contained by a representation, the place to look is not to the extensional

notion of what is contained within the entity represented, but rather to some kind of intensional notion.

In light of these considerations, a better proposal, I believe, appeals to the notion of *entertainment*. This is an intuitive notion in widespread use in philosophy, and it is not easy to define in more precise terms. Very roughly, S entertains p iff p is in an appropriate sense 'before her mind'. The advantage of appealing to entertainment is that it is a reasonably well-understood, or at least common, way in which the content of a propositional attitude can be 'detached' from what it entails. For instance, a person who entertains that everything that exists has a color need not, it is very reasonable to think, entertain that everything that exists has a shape, even if, as is plausible, everything that has a color has a shape. This is promising because what we are looking for is a way to distinguish propositions that are contained by a subject's overall global representation from propositions that are not.

At a first pass, we might try to use the notion of entertainment to refine the previous attempts at accounting for what it is for a subject to have an overall representation that contains everything else she believes as follows:

A subject S has an overall global representation iff there is a proposition p such that S believes p and $\forall x$ (x is a proposition that S believes $\rightarrow$ p entails x and, necessarily, whoever entertains p entertains x)).

In this formulation, which is inspired by Chisholm,³ the relevant notion of containment is captured by the idea that p entails every proposition the subject believes and is such that, necessarily, entertaining p entails entertaining every proposition the subject believes.⁴ The model has three components. First, like the earlier models, it understands the unity of belief to consist in a subject's believing some single proposition. Second, in believing that

³ See Chisholm 1986, p. 73.

⁴This should not be taken to imply that in order to believe p , the subject must be able to entertain p .

proposition the subject believes a proposition that is true only if everything she believes is true. Third, there is a condition to rule out opaque entailments – it is not just that the single proposition p is true only if everything the subject believes is true, it is that each thing the subject believes is itself represented by or within p .

This formulation is an improvement, I believe, but a further refinement is needed. One might think that there could be some proposition q the subject does not believe, but which is nonetheless entailed by p and such that whoever entertain p entertains q . This would mean that whether the overall global representation that has p as content has q as part of its content depends on whether the subject believes q , and this cannot be right. (The objection is by now familiar.) As an example of the kind of case I have in mind, suppose p is the conjunctive proposition:

everything that exists has a color and (everything that exists has a color or everything that exists has a shape)

This proposition entails that everything that exists has a shape (on the assumption of the metaphysical principle that everything that has a color has a shape), and is such that whoever entertains it necessarily entertains ‘everything that exists has a shape’. So if the model most recently considered were correct, then if there were some subject who believes that proposition p , and that everything that exists has a shape, she would have an overall global representation just by virtue of believing p . Moreover, p would be the content of their overall global representation, and thereby contain that everything that exists has a shape. This cannot be right.

What we need, therefore, is

A subject S has an overall global representation of the world iff there is a proposition p such that S believes p and $\forall x$ (x is a proposition that S believes iff p entails x and, necessarily, whoever entertains p entertains x)).

This refinement would seem to rule out this class of objection.

If this general account is correct, then it also allows us to give a theoretical account of the relation between a subject's individual beliefs and her overall global representation. After all, it presumably is not merely that a subject has an overall representation that has as its content everything that is the content of her individual beliefs, it is that those individual beliefs in some sense compose the overall global representation. There is, that is, something akin to a mereological or compositional relationship between her individual beliefs and her overall global representation. The general account given above lends itself to an analysis in which the subject's individual beliefs are *doxastic parts* of her overall global representation:

A belief x with content Cx is a doxastic part of a belief y with content Cy iff (i) Cy entails Cx , (ii) necessarily, whosoever entertains Cy entertains Cx , and (iii) necessarily, whosoever has y has x .

Using this definition of 'doxastic part', it is possible to give an alternative rendering of the general account already on offer:

S has an overall global representation of the world = S has a belief B such that $\forall x$ ($(x$ is a belief and S has x) iff x is a doxastic part of B)

This is, I take it, simply the analog of the formulation given above recast in terms of beliefs.

Conjunctive Belief and the General Model

This proposed formulation specifies a general condition a subject has to meet if she is to have an overall global representation, without saying how, exactly, she is to meet that condition. However, it is worth pointing out that this account appears to fit very well with what is perhaps the most natural option for making sense of global representation. An attractive idea (attractive, anyway, to a lot of other people, a qualification the point of which

will become clear shortly) is that the content of an overall global representation, and the having of such a representation, is to be understood via appeal to the notion of conjunction. On this approach, the content of a global representation is a conjunctive proposition, and a subject has an overall global representation if and only if she believes whatever proposition is the conjunction of the propositions that are the content of what we usually call her individual beliefs. This is, I think, almost certainly too strong, for even if it is true that a subject has an overall global representation if she believes such a conjunctive proposition, it seems that it is at least metaphysically possible for some believer to have an overall global representation other than via conjunction. Hence the conjunctive proposal should, by those who favor it, be thought of as complementing rather than as replacing the general account offered above: believing a conjunctive proposition would be one of the ways a subject could meet the condition offered by that account (leaving open that there may be other ways of doing so). For suppose the subject's 'individual beliefs' were that p_1 , that p_2 , that p_3 , ..., that p_n , and the subject believed the proposition that p_1 and p_2 and p_3 and ... p_n . This latter proposition, which we can call the subject's global conjunctive proposition, would (i) entail each proposition the subject believes, since any conjunction entails each of its conjuncts, (ii) be such that, necessarily, anyone who entertained it would entertain each of p_1 ... p_n , and (iii) be such that, necessarily, anyone who believed it believed each of p_1 ... p_n .

If one adopts both aspects of the story given so far – the general account proposed, and the conjunctive model that is the perhaps the most natural account of global representation – one could still go in two quite different directions, depending on the conditions under which one takes a subject to believe her global conjunctive proposition. Many philosophers would set those conditions such that few, and probably no, actual human subjects believe their global conjunctive proposition – and hence such that, at least by appeal

to conjunctive belief, few and probably no actual human subjects have an overall global representation. There are three kinds of considerations that one might think weigh against a subject believing her global conjunctive proposition. First, one might think that to believe a conjunction of any kind requires, or at least typically requires, some kind of deliberate, perhaps even conscious, effort or activity on the part of the believing subject. On this picture of what is involved in believing a conjunctive proposition, while some of us are better and others worse at ‘conjoining’, as the mental activity may be called, it is a safe bet that none of us ever have gone about this activity with the plodding dedication that would be required were we to achieve belief in our global conjunctive proposition. Second, one might maintain that a subject can believe a proposition only if she has the capacity to consciously entertain the proposition, and insist that a subject could not entertain her global conjunctive proposition. Richard Foley, for instance, makes this claim in a footnote, saying:

If persons have a very large number of beliefs, [the claim that a person believes her global conjunctive proposition] can be true only if persons have incredibly complex beliefs. But in order to believe something $p \dots$ a person must be psychologically able to consider p . Yet, it seems very unlikely that any person is ever able to consider the conjunction of the contents of all his beliefs. Thus, [the claim] seems to be false. (1979, p. 249, fn. 7)

Third, one might claim that beliefs are constrained by consistency, and that this norm would be violated were a subject (as least if she is anything like actual human subjects) to believe her global conjunctive proposition. For instance, one might maintain that it is impossible to believe a contradiction, and claim that this rules out that a subject (again, if she is anything like a typical human subject) believes her global conjunctive proposition. In effect, this is what Stalnaker, Lewis and Davidson do by assigning inconsistent beliefs to separate belief compartments.⁵ Alternatively, one could reject all three lines of thought and take it that a

⁵ One might think the general model itself prevents anyone with inconsistent beliefs from having an overall global representation, on the grounds that it requires the subject to believe a proposition that entails every proposition the subject believes, and no proposition entails propositions that are inconsistent. But this

subject does, at least typically, believe the conjunction of the propositions that are the content of what we call her individual beliefs. Simon Evnine appears to endorse this position in his 1999 paper “Believing Conjunctions”, offering the following considerations in its defense:

I begin by giving an argument for (Conj) [the claim that if a subject believes that p and believes that q , the subject believes that p and q]. The argument is that, in normal circumstances, its truth is presupposed by our practice of attributing conjunctive beliefs. Take any two propositions that are of no special logical or emotional significance to someone, say that grass is green and that snow is white. I contend (this is intended as an empirical observation about our practice) that if we have good grounds for attributing belief in each of these to a person, then, absent any special circumstances, that is all we need to attribute to that person a belief in their conjunction. Or if, for example, we are summing up a position someone has just explained at some length, we can do so by attributing to that person a large conjunctive belief. This conjunctive belief describes in a single proposition a number of different propositions that were expressed severally over a period of time. We do not require evidence that, for example, some concrete psychological process has occurred, a process in which the two individual beliefs are put together or conjoined. Actual belief in the conjunction of one’s beliefs is the default mode in belief attribution. (1999, pp. 214-215)

These two different directions of response agree on the two-part model offered above, and simply disagree about when a subject meets those conditions. Anyone who thinks of conjunctive belief as the product of a deliberate or conscious activity of conjoining would surely say that as a matter of empirical fact, no one believes his or her global conjunctive proposition. Foley thinks human subjects are constitutively barred from believing their global conjunctive propositions. Stalnaker, Lewis and Davidson think human subjects seldom believe their global conjunctive proposition. And someone like Evnine thinks human

objection does not seem right, for, standardly, a necessarily false proposition (such as one containing a contradiction) entails every proposition, and hence inconsistent propositions. So long as a subject can believe a necessarily falsehood, therefore, a subject could believe a proposition that entails inconsistent propositions. What seems clear is that if a subject cannot believe a necessarily falsehood, then the general model would bar any subject with inconsistent beliefs from having an overall global representation. Some may think this a virtue of the model, others may think it a vice.

subjects at least typically do. Any is in a position to say: “Great, we have an account of global representation. End of Story.”

Concerns with Conjunction

The appeal to conjunctive belief appears to offer a tidy account of global representation, and one that fits well with the more general model. However, there may be reason to think the appeal to conjunctive belief is not correct. I am not able to articulate my concerns very well, or very fully, and I have reached no firm conclusion on the matter. I think it is important, however, to do what I can to register and develop these worries, even if I am not at this point satisfied with what I have to say on the issue.

The first reason I have doubts about the conjunctive account is that I do not understand what a conjunctive belief is. A second reason is that believing a conjunctive proposition seems not to capture the kind of unity that is involved in having a global representation. A third reason is that, whatever we make of conjunctive belief, the possibility of having such a belief, and indeed of being a believer in any intelligible sense, rests on another kind of doxastic state, and this state (as I will argue in the next chapter) is a better candidate to explain the unity of belief.

The balance of this chapter will be devoted to developing these doubts about the correctness of the conjunctive account of global representation.

Concerning the first worry, that it is unclear what a conjunctive belief is: The usual assumption is that it is possible for a subject to believe that p , and to believe that q , without believing the conjunctive proposition that p and q . Such a subject is said to lack the conjunctive belief that p and q by virtue of having failed to “put the beliefs together”. However, it is a widely shared – and deeply grounded, in my view – assumption that

conjunctive facts or states of affairs are nothing other than their conjuncts. It is not just that they supervene on their conjuncts; it is that they are no addition to being. It seems to me that these two assumptions are in an unhappy tension. For if the first assumption is true, then is possible for a subject to believe that p and that q , but nonetheless fail to believe that p and q , in the sense that she fails to represent the world as being such that p and q . Yet if the second assumption is true, then it is hard to see how this conjunctive fact escapes her ken, given that it is nothing other than facts (that p and that q) that she already represents as being the case. As Wittgenstein put it:

4.26 If all true elementary propositions are given, the result is a complete description of the world.

It seems to me that if I believe every true non-conjunctive proposition, my representation of the world would be perfect. If one objects and says – No, you would be missing out on countless conjunctive facts! – I can only say in response that I would not be missing out on these facts since these facts are nothing more than facts I am aware of. My world, if the expression be permitted, would not change for coming to believe the conjunctive facts.

We are familiar, of course, with ways in which it seems a person can be ignorant of some fact, even though it is hard to understand how the fact they seem to be ignorant of is anything other than some fact they already know. Two kinds of cases are prominent, but both involve considerations that seem distant from what is involved in believing conjunctions. First, think of Lois Lane, who appears to be ignorant of the fact that Clark Kent can fly, even though that fact is nothing other than a fact we are inclined to admit she knows, which is that Superman can fly. But here there is a familiar opacity that rests, somehow, on the fact that Lois Lane has two different ways of thinking about one entity. It is hard to see how any such story could apply in the case of conjunction. Second, think of

the relationship between the facts expressed by the vocabulary of the special sciences and the facts expressed by the vocabulary of physics. The fact that the continental plates are drifting at such and such a speed is not (I think and it is plausible to maintain) anything other than that such and such is going on at the microphysical level. Yet, one wants to insist, a subject can surely be ignorant of the first (special science) fact without being ignorant of the second (physics) fact. It is not easy to sort out what is going on in this kind of case, to put it mildly, but the problem seems to be rooted in the fact that the special science 'level' and the physics 'level' involve concepts that are utterly different. Again, nothing at all like this is going on with conjunctive propositions. What is needed, therefore, is a clearer understanding of how the state of believing both that p and that q differs from the state of believing that p and q , given that the conjunctive fact that p and q is nothing other than its conjuncts and that conjunctive facts do not seem to involve any kind of conceptual opacity, which is the only way we currently have, I take it, of understanding how a person could fail to believe something even though it is, extensionally, equivalent to something she already believes. Or, to put the point another way, we have to understand how the totality of true non-conjunctive propositions *do not* completely describe the world.

These complaints barely scratch the surface of what I think are difficult issues. Moreover, it is not possible, here, to dig much deeper – and the bottom is a long way down. My goal here is not to refute the conjunctive model of global representation, however; it is only to motivate some worries. At a minimum, the consideration raised suggests that we need a much better understanding of conjunctive belief than we now have: we need to understand, more clearly, what exactly someone who believes every non-conjunctive proposition is missing out on. It is not enough, I think, to say simply that such a person is ignorant of a conjunctive fact, for we have to understand how this is possible, or what this

amounts to, given that she is not ignorant of the non-conjunctive facts that the conjunctive fact in some senses reduces to.

Of course, it might be that one should accept that $B(p \ \& \ q) \neq (Bp \text{ and } Bq)$ even if one does not understand how the claim could be true. As Thomas Nagel has pointed out in another context, it is possible to have sufficient evidence to believe a proposition even if one does not understand how the proposition could be true. (1974) However, I do not think standard arguments for the distinction are convincing. There is not room here to consider the issue in detail, but I will briefly indicate where I think standard arguments go wrong.

Belief and Entertainment

First, in regard to the principle that one cannot believe a proposition unless one is ‘psychologically able’ to consider it. This can be taken to support a distinction between $B(p \ \& \ q)$ and $(Bp \ \& \ Bq)$ so long as one thinks there are situations in which a person can believe propositions $p_1, p_2, \dots, p_n$ without being able to entertain the relevant conjunctive proposition. In the short passage quoted from Foley, he offers no defense for this principle. He is not alone in thinking it needs no defense; it often appears as an undefended premise in an argument, with the apparent assumption that it is either self-evident or uncontroversially true. For instance, in a recent *Philosophical Studies* paper Mark Eli Kalderon rests an argument upon the principle, but relegates the assertion of it to a parenthetical remark, as if it is too obvious to need to put in the normal part of the sentence let alone to defend.⁶ A second example occurs in an October 2007 article in *Analysis*, in which Krister Bykvist and Anandi

⁶ “In order to know the proposition expressed by the standard explanation [of what it is to function semantically as an existential quantifier] the speaker must already be able to entertain the content of existential quantification (since one can only know propositions that one can entertain).” (2001, p. 141.)

Hattiangadi repeatedly appeal to the principle⁷ (sometimes in the context of whether one can believe a ‘long’ conjunctive proposition) without offering a word in its defense. This is so even though the general rigor of the piece makes it clear that they take themselves to be in the business of giving serious arguments (unlike perhaps Foley, who compressed his remark into a footnote). In my view, the principle is often asserted without defense because it is conflated with the much more plausible claim that one cannot believe a proposition unless one has the conceptual resources to understand the proposition. Much of the intuitive motivation for the idea that one can believe a proposition only if one can entertain it derives, I suggest, from the fact that usually when we imagine a person who cannot entertain a proposition, we imagine a case in which they cannot entertain it because they lack the necessary concepts. For instance, Aristotle didn’t have the capacity to entertain the thought that igloos are best built with the door facing away from the wind, nor did he have the capacity to believe it. But it isn’t that he couldn’t believe it because he couldn’t entertain it – the lack of each capacity had a common cause, or determination base – his lacking the concept ‘igloo’.

To illustrate the distinction, and why it is easy to overlook, I’ll examine two arguments offered by Robert Audi, both directed toward the conclusion that people do not have an infinite number of beliefs (this issue will not concern us). In the first argument, offered in a 1982 paper, Audi is careful to distinguish between being able to entertain a proposition and having the conceptual resources to understand the proposition. But writing

⁷ For instance, “take the conjunction of all the necessary truths – a proposition that is far too complex for you to grasp. According to (2b) [(2b) For any S, p : S ought to (believe that p if and only if p is true)], you ought to either bring it about that p is false or bring it about that you believe that p . But you can do neither. You cannot bring it about that p is false because p is a necessary truth. And you cannot come to believe that p because it is not humanly possible to grasp such a complex proposition, let alone believe it. Since ‘ought’ implies ‘can,’ (2b) must be false.” (2007, p. 279)

some 15 years later in an introduction to epistemology, and offering what appears to be much the same argument, Audi is less careful. I will argue that the distinction is genuine and that it is a mistake to think that one must be able to entertain a proposition in order to believe it.

Audi's first argument:

[T]here remains a major obstacle to the view that one might have infinitely many mathematical beliefs. Consider the generally accepted point that S can believe that p only if S can grasp that p , in a sense that entails understanding it. Perhaps he need not be able to entertain it all at once, but he must at least be able to 'take it in', as by reading a formulation of it comprehendingly. Now consider the infinite series 1, $1\frac{1}{2}$, $1\frac{3}{4}$, $1\frac{7}{8}$,....Expressing each of these elements requires infinitely many numerals (or infinitely many symbols). Some of the elements, say those whose expression takes over a million numerals, will simply be too cumbersome for S to grasp in the relevant sense. If in some sense S grasps them all by virtue of comprehending that 2 is greater than any number in the series, this does not imply his grasping each in such a manner that he may be said to believe that 2 is greater than nWhereof one cannot grasp, thereof one cannot believe. (1982, p. 118)

Audi begins by noting that it is generally accepted that one has to be able to understand a proposition in order to believe it, which he takes to involve being able to "take it in", as by "reading a formulation of it comprehendingly". He distinguishes being able to do this from being able to "entertain it all at once", which he grants may not be necessary. Neither expression is altogether clear, but I cannot do much to improve on Audi's wording. He then offers an argument concerning whether a subject can have an infinite number of beliefs of a certain form that makes it clear he thinks they cannot because the subject cannot 'take in' the proposition, and not because the subject fails to be able to entertain the proposition 'all at once'. Roughly, the central idea is this:

A subject cannot have an infinite number of beliefs of the form '2 is greater than 1', '2 is greater than $1\frac{1}{2}$ ', '2 is greater than $1\frac{3}{4}$ ', '2 is greater than $1\frac{7}{8}$ ', etc., for eventually, by following this schema, we will arrive at a proposition that, when expressed in English, is such that the second number named is written using too

many numerals to be grasped. For instance, this will be so at least for all of those whose proper name involves more than one million numerals.⁸

What Audi is imagining here is that, by following the schema, we will come to some proposition that would be written in English as ‘2 is larger than n ’, where ‘ n ’ should be read in this formulation as standing in for some very long string of numerals. And his point is that since we cannot grasp any item in the infinite series $1, 1\frac{1}{2}, 1\frac{3}{4}, \dots$ whose expression takes over a millions numerals, we could not understand and hence could not believe a proposition of the form ‘2 is larger than n ’ that has such an element as a constituent.

However, it is clear that in this case the force for thinking we couldn’t believe it comes from the fact we fail to grasp one of the proposition’s constituents, not from the fact we cannot entertain the proposition. In other words, the limits that bar belief are *conceptual*. In the case imagined the subject can neither entertain the proposition ‘all at once’ or even ‘take it in’, but what rules out that he believes it is that he cannot even ‘take it in’, for he has no conceptual access to the number to which the number 2 is asserted to stand in the greater than relation. (I should also note that in this case the proposition itself is rather simple, for it does no more than assert a relation of size between two entities. Whatever it might mean for a proposition to be ‘long’ and ‘complex’, this one is neither.)

But now consider how Audi presents a superficially similar argument in *Epistemology*:

A Contemporary Introduction, a 1998 textbook:

It is doubtful that, given our psychological make-up, we *can* know, or even believe, infinitely many things. It might seem that we can have an infinite set of arithmetical beliefs, say that 2 is larger than 1, that 3 is larger than 2, and so forth. But surely for a finite mind there will be some point or other at which the relevant proposition cannot be grasped (the point might be different for different people or even the same person at different times). Imagine the “largest” proposition a supercomputer

⁸ It doesn’t matter if we think of the relation between the term ‘ n ’ and the number n as naming or some other relation; the point is just that however that element of the proposition is expressed in an English sentence, the term that expresses or names it will be too cumbersome to grasp at least whenever the term is one million numerals long.

could formulate after years of work. It could easily be too long to understand and so cumbersome that one could not even take in a formulation of it. One would be unable to remember enough about the first part of it when one gets to the end; one could thus never understand the whole thing. What we cannot understand we cannot believe; and what we cannot believe we cannot know. (1998, p. 189)

The argument begins much the same way, but Audi is less clear about the fact that in the case imagined the proposition cannot be grasped because the would-be believer has no cognitive access to one of the proposition's constituents. The problem instead is presented as that of the proposition itself being too 'long' and cumbersome. But this isn't right: the proposition is no 'longer' or more cumbersome than that expressed by the sentence '2 is larger than 1'. In any event, the pair of arguments suggests how easy it is to slide between the claim that if the English term for an entity is so long that one cannot 'take it in', one cannot (at least by reading the term) form a concept of the entity and thus cannot understand and hence believe a proposition concerning it, to the claim that if a proposition is too 'long' to entertain, in the sense of having the English sentence expressing it run before one's mind, then one does not understand the proposition and hence cannot believe it. The distinction is relevant because in the case of 'long' conjunctive propositions, whenever one believes each conjunct, one does not fail to have the conceptual resources to understand the proposition. A person cannot, we can grant, 'hold it all before her conscious mind at once', but that is not clearly relevant.

One might think, however, that the points Audi raises are enough to suggest obvious support for the claim that to believe a proposition one must be able to entertain it. The key idea is that to believe a proposition one has to understand it. The argument could then say that although in Audi's cases the inability to understand the proposition rests on a person's inability to grasp one of the elements of the proposition, the general point that one has to understand a proposition in order to believe it rules out that a person can believe a

proposition they cannot entertain. Then, on the assumption that one cannot entertain one's global conjunctive proposition, it follows that one does not believe it. We can put the argument as follows:

Necessarily, (for all S and p) if S believes p , S understands p .

Necessarily, if S understands p , S can entertain p .

Necessarily, if S believes p , S can entertain p .

Once the argument is put plainly, however, it is far from clear that it is sound. The reasons for my thinking so can be found within the passage from Audi above. I think Audi got it right the first time, in the earlier paper, when he said:

Consider the generally accepted point that S can believe that p only if S can grasp that p , in a sense that entails understanding it. Perhaps he need not be able to entertain it all at once, but he must at least be able to 'take it in', as by reading a formulation of it comprehendingly.

If we distinguish between understanding a proposition and being able to entertain it, then premise 2 suddenly requires a defense; it is not trivially true. Moreover, I think there is good reason to think premise 2 is false.

We can see that being able to understand a proposition does not entail being able to believe it, and hence that premise 2 is false, but considering a thought experiment. Think of someone you know well, whose behavior is such that you would not hesitate to ascribe to her the belief that (say) Jaegwon Kim is not identical to Ernest Sosa. But now suppose that some doctor discovered that the person in question had the following psychological, or neurophysiological, condition: her brain is wired such that she cannot entertain that proposition – it just so happens that there are certain causal mechanisms in her brain that prevent her from being able to do so, or trigger massive interrupters were she ever on the verge of entertaining the proposition. I do not think this would give us compelling reason to think that, therefore, she fails to understand the proposition that Jaegwon Kim is not

identical to Ernest Sosa. Imagine, for instance, that the person in question is a secretary in the philosophy department at Brown University. She grasps the concept ‘Jaegwon Kim’,⁹ the concept ‘Ernest Sosa’, and the relation of non-identity. Moreover, when she fields telephone calls, sorts mail, greets faculty in the morning, and so on, she shows no confusion concerning who is who. She just never has, never will, and cannot frame or entertain the thought that Jaegwon Kim is not identical to Ernest Sosa. The most reasonable interpretation of this case, I think, is one in which the secretary does understand the proposition – she just cannot entertain it.

Note, though, that if one resists this, and insists that the secretary does not understand the proposition in question, then what has to fall is the first premise of the argument, the one that says a subject can believe a proposition only if they understand it. For surely in the case imagined the secretary does believe that Jaegwon Kim is not identical to Ernest Sosa. Otherwise her competence in sorting mail, directing telephone calls, saying hello in the morning, etc., seems inexplicable.¹⁰

Finally, here is an additional reason to doubt the truth of premise 2. Consider the following proposition:

1 is a number and 2 is a number and100,000,000 is a number.

This proposition has 100 million conjuncts. If ‘entertain’ means ‘run through the English expression in one’s head’, then I cannot entertain this proposition. But it seems to me that I

⁹ Or is a competent user of the name ‘Jaegwon Kim’, if you prefer.

¹⁰ One might object here that she has the capacity, but there is merely an interrupter of sorts, in the sense that absent the interrupter she would be able to entertain the proposition. And, the objection continues, this is different from the case of considering a long conjunction proposition, since one isn’t psychologically capable of entertaining that for reasons that have nothing to do with an interrupter. I don’t think this objection succeeds. In the case of entertaining a conjunctive proposition, there is an interrupter of sorts – one’s attention span, chiefly, the length of one’s life, one’s flagging memory of the earlier parts of the proposition, and so on. The constraints on one’s entertaining a conjunctive proposition seem to be constraints that are relevant to entertaining it, and not constraints that are relevant to believing it.

do understand the proposition, and moreover that I believe it. Foley, and perhaps Audi, would be committed to saying not only that I don't understand it or believe it, but that it is at least psychologically impossible for me to believe it. This seems to me to get the facts wrong. My grasp of each of 1, 2, 3, 4....100,000,000 is strong enough, I think, that I grasp what this conjunctive proposition asserts and believe it.¹¹

I have suggested that much of the intuitive force behind the idea that a person can believe a proposition only if she can entertain it rests on conflating 'S cannot entertain that p' with something like 'S lacks the concepts necessary to entertain that p'. Moreover, I think consideration of a couple of examples makes it very reasonable to think that a person can believe a proposition without being able to entertain it. There is, however, a functionalist line of argument in support of the claim that belief requires the capacity to entertain which I should also consider (I am indebted to Chris Hill for bringing this line of argument to my attention). Doing so will, I hope, show that functionalist considerations actually support, rather than weigh against, the claim that a subject can believe a proposition even if she cannot entertain it.

The functionalist objection runs roughly as follows:

Belief states are defined functionally, that is, by appeal to their normal causes and effects. One of the effects of a belief that p is that one has a tendency to assert that p, or to use that p in an overt piece of reasoning, where to do either of these things requires that one be able to entertain that p. Therefore, believing a proposition entails being able to entertain the proposition.

The idea behind the objection is that to believe a proposition is to be such that, among other things, one is disposed to behave in ways that involve entertaining that proposition. Thus, to believe that p seems to require that one be capable of entertaining the proposition in

¹¹ I am not arguing that this would continue with ever higher numbers; it may be, for instance, that I just don't have a clear enough grasp of numbers beyond 100 million. We need not decide the issue.

question. After all, how could one be disposed to behave in ways that involve entertaining the proposition if one is not capable of entertaining the proposition?

I think this argument is interesting but ultimately unsuccessful. Let us grant functionalism.¹² What we have to ask is whether the following is true: Necessarily, for all *p*, if *S* believes *p*, *S* is disposed to act in ways that entail that *S* entertains *p*. I don't see reason to think this premise is true. Functionalism requires only that belief states be defined functionally; it doesn't specify that the causal role of a belief, necessarily, involves disposing the believer to behave in ways that involve entertaining the proposition. We should probably grant that 'ordinary' beliefs, such as the belief that the cat is on the mat, or that it is going to rain sometime this week, or that pancreatic cancer is more deadly than liver cancer, etc., at least typically dispose those who have them to behave in ways that involve entertaining the proposition. But it is a big step from here to the claim that *all* beliefs *necessarily* dispose those who have them to behave in ways that involve entertaining the relevant propositions. We have good reason to think, I believe, that there could be cases in which a subject has even an 'ordinary' belief, and yet is not disposed to act in ways that entertain it. (The secretary in the Brown philosophy department, imagined above, believes that $JK \neq ES$ but is not disposed to entertain that proposition. The belief is still functionally grounded in her behavior, however: she greets them differentially, sorts mail successfully, etc.). Moreover, even if having the belief that *p* necessarily disposed a subject to entertain that *p*, where *p* is an 'ordinary' proposition (roughly, of a kind that can be expressed in a reasonably compact natural language sentence), this wouldn't show that believing any proposition necessarily disposes a

¹²The argument I am here considering could be put in terms of a reductive functionalism, on which belief states simply are functional states, or in terms of a view of belief on which belief states, or at least belief state attributions, must be grounded in behavior of some kind (where this need not be because they are identical to functional states). I don't think for the purposes of considering this argument we need to distinguish carefully between these two versions of the argument, for they would agree on the relevant point, that a belief that *p* is necessarily connected to certain kinds of behavior that involve the entertaining of *p*.

subject to behave in ways that involve entertaining it. After all, it is quite easy to entertain the proposition that the cat is on the mat, and quite difficult to entertain a many-place conjunctive proposition. Why should we expect that if believing the former entails being able to entertain it, then believing the latter entails being able to entertain it as well? What we would need is clearer reason to think that something is a belief only if it shares that functional role with 'ordinary' beliefs. This is not an unreasonable view, especially since what I am calling ordinary beliefs are paradigmatic beliefs, and those beliefs through which we have our central grasp on the concept 'belief'. But we should remember that beliefs are 'anchored' in observables through much more than behavior that involves the entertaining of the believed proposition. Consider again the Brown secretary: A proponent of the functional argument I am considering would seem committed to holding that functionalism shows that she does not believe that $JK \neq ES$, whereas functionalism seems, in fact, to show the opposite. If, in general, a person alleges that S believes p, functionalism requires that S's belief that p be in some important way grounded in her behavior. This could be through her being disposed to assert that p or entertain that p, but it could be grounded just as firmly, I think, in other kinds of behavior. For these reasons I think that commitment to a functionalist account of belief at best generates a warning that if a belief does not dispose a subject to engage (in appropriate circumstances) in behavior that involves entertaining the proposition believed, the belief must be in some other way functionally anchored.

If these considerations are correct, then if there is a difference between $B(p \ \& \ q)$ and $(Bp \ \& \ Bq)$, this needs to be shown by appeal to something other than the principle that one can believe a proposition only if one can entertain it, coupled with the premise that one cannot entertain conjunctive propositions that have as conjuncts the non-conjunctive propositions that one believes.

Belief and Consistency

A second line of thought that one might think supports a distinction between $B(p \ \& \ q)$ and $(Bp \ \& \ Bq)$ concerns consistency and contradiction. The worry rests on the fact that human subjects are less than perfectly rational, and insists that their irrationality deepens, perhaps to the point of being unintelligible, if it is the case that a subject believes that p and q whenever she both believes that p and believes that q . Here are two different reasons one might think so.

First, one might hold that propositional attitude ascriptions are bound by norms of rationality, and that these norms would be violated if there were no difference between believing a conjunctive proposition and being such that one believes each conjunct. For example, one might think, there are situations such as the following: a subject has the desire to Φ whenever both p and q , is such that she both believes that p and believes that q , but nonetheless fails to intend to Φ (in a way that isn't explained by an all-things-considered judgment).¹³ One way to try to make sense of her behavior is to say that she has 'failed to put the two beliefs together' – she is rationally remiss, but not as grievously as she would be were she to believe the conjunctive proposition that p and q but nonetheless fail to act. Moreover, if, as Davidson and Lewis emphasize, propositional attitudes are identified by their place within a network (as Davidson puts it), or constitutive rationality determines the

¹³ Perhaps an example of this kind of case would be as follows: Gurpreet has heard that two acquaintances of his, Selena and Yomna, do a great comedy routine when they get together at parties. Gurpreet has, in general, little interest in parties, and generally no desire to attend them. He does, however, form the desire to attend the next party he hears about at which he expects to find both Selena and Yomna, as he figures that the comedy routine should be worth the discomfort and trouble of attending. A few weeks later he runs into Yomna, who mentions an upcoming party she plans to attend. Later, he hears from Selena that she is going to the same party. At that point he is in a state such that he believes that Selena is going to the party and believes that Yomna is going to the party, but he hasn't formed any intention to go.

content of propositional attitude (as Lewis puts it), then it becomes very hard to make sense of that kind of attitude attribution. It just is not possible, the worry insists, to desire to Φ whenever both p and q , believe at t that both p and q , yet fail at t to intend to Φ (unless, again, one refrains from intending because of an all-things-considered judgment).¹⁴

Second, one might worry that if there is no difference between being such that one both believes that p and believes that q , and being such that one believes that p and q , then some subjects, perhaps even many human subjects, believe contradictions. And this, one might think, is just impossible. The idea here is that while it is possible for a subject to believe (at t) both that p and that $\sim p$, it is not possible for a subject to believe that p and $\sim p$. Thus there must be a distinction between believing each of two propositions and believing their conjunction.

These are deep and important problems that I cannot address fully here. I can, however, say a few words to indicate why I do not find them convincing grounds for thinking there is a difference between $B(p \ \& \ q)$ and $(Bp \ \& \ Bq)$. What I will have to say will not give anyone (including me) reason to believe these lines of argument *cannot* be used to support the distinction, but it may give readers (and at least gives me) reason to think they do not clearly trump the general worry I have about the intelligibility of the claim that $B(p \ \& \ q) \neq (Bp \ \& \ Bq)$.

The basic problem I have with both arguments is the same, but there is an additional point to make concerning each that I will raise first to get out of the way. Concerning the first line of thought, it seems to me that in cases of the kind imagined, we can always say the behavior (understood in the most general sense to include the formation and maintenance of

¹⁴ I should note that Lewis does not claim something this strong. He argues that *types* of propositional attitudes are governed by ‘constitutive rationality’, and wants to leave open that occasionally tokens could fail to meet those conditions. See 1986, p. 39.

propositional attitudes) shows not that the subject fails to believe the conjunctive proposition, but that the belief fails to connect in an appropriate way with behavior. It is not that the subject has not put the two beliefs together, whatever that means, it is that the subject's believing fails to guide her action as it should. If the subject in question desires to Φ whenever both p and q , believes that p and believes that q , and nonetheless fails to intend to Φ , the problem may be that her beliefs and desires are not causing what they ought to. We do not need to imagine that the problem concerns the content of her believing, rather than how her believing is connected through desire to action and intention. Concerning the second line of thought, there are two things to say. First, it is not entirely clear that a subject cannot believe a contradiction. Evnine (in "Believing conjunctions") argues subjects do believe contradictions, but I am content to say the case has not been convincingly made that they cannot. Stalnaker says that "one cannot agree to disagree with oneself" (1984, p. 84), but doesn't make clear how believing a contradiction would involve doing this. And Davidson, while giving us reason to believe that inconsistency couldn't be the norm within a belief system, does not really show that it is impossible for a subject to believe some contradiction.¹⁵ If one is committed to certain kinds of accounts of the semantic content of belief states, such as a possible worlds model, one could not accept that a subject believes a contradiction, at least without rejigging the theory substantially. (For instance, one might attempt to develop a model that takes the content of *atomic* beliefs to be given by sets of possible worlds, and takes conjunction, disjunction, and other logical connectives to be formally defined operations over or between such contents. I do not know if any model like this has been explored.) But that is not the same as arguing for the claim. And, finally, I do

¹⁵See Davidson 2004h, especially page 217, for a defense of the claim that although a subject can believe contradictory propositions, she cannot believe a contradiction.

not find convincing the worry that it just doesn't make sense to imagine that the semantic content of a state is inconsistent. For maps, dictionaries, paintings and so on can contain contradictory propositional content. Consider for instance a map of the earth, in Mercator projection, with a scale indicating distance and countries shaded progressively from dark to light based on increasing total land area. Because of the scale distortion in higher latitudes with a Mercator projection, the map presents (via the scale metric) Greenland as being many times larger than the Sudan.¹⁶ Yet the map also presents (via the shading metric) the Sudan as being larger than Greenland (it is in fact 115% the size of Greenland). The map thus contradicts itself – it says that Greenland is many times the size of the Sudan, and that that the Sudan is larger than Greenland. I see no barrier to thinking, however, of the content of the map as being a proposition – one that could not possibly be true. Therefore, it doesn't seem that one can argue that it is impossible to believe a contradiction on the grounds that nothing could have a necessarily false proposition as its content. Second, the objection or argument being considered insists that it is possible to believe that p and that $\sim p$, but impossible to believe that p and $\sim p$. I have thus far expressed skepticism about whether it is in fact impossible to believe a contradiction. But one might equally accept this is impossible and worry, instead, about whether it is really possible to believe both that p and that $\sim p$. In what situation, for instance, could such a belief state ever come to light?¹⁷

The more basic problem I have with both lines of thought, however, is that they critique the idea that a subject always (or automatically) believes a conjunction whenever she believes each conjunct, in the face of *prima facie* problems, by proposing alternative ways of

¹⁶ Mercator projections represent Greenland and Africa as being roughly the same size, when in fact Africa is nearly 14 times larger.

¹⁷ See Wedin 2004 for a very interesting discussion of Aristotle's view of conjunctive belief. Aristotle apparently held both that it was not possible to believe a contradiction, and that there is no difference between believing a contradictory proposition and having contradictory beliefs.

understanding the subject's mental states that are solutions in name only. I take seriously the kind of consistency constraints pressed by Davidson and Lewis, and am sympathetic to an unstructured (e.g., possible worlds) model of belief content. And I recognize that these constraints run up against some puzzling features of the mental, such as that we appear to have beliefs with inconsistent contents, or that our propositional attitudes are not fully rational. What the lines of thought that seek to establish that there is a difference between believing that p and that q and believing that p and q have in common is that they want to allow for inconsistency while also shunning it. Hence the appeal to compartmentalization, for instance. That move is motivated by, and in fact rests upon, accepting the kind of constraints that suggest there is no difference between believing that p and q and believing that p and that q . The problem is that we just have no idea what it is to have belief compartments. It may not even be much of an exaggeration to say that our conceptual grip on the concept 'belief compartment' just is something like 'that theoretical apparatus that will allow us to say beliefs (or propositional attitudes in general) must be consistent while allowing that a subject's beliefs (or propositional attitudes) don't have to be consistent'. For all the good concern with consistency, this amounts to having our cake and eating it too.¹⁸

A third and final line of argument that one might think establishes that $B(p \ \& \ q) \neq (Bp \ \& \ Bq)$ rests on epistemic considerations, specifically on the apparent fact that one's justification for believing a conjunctive proposition is not the same as one's justification for believing each conjunct. We can see the nature of that issue, and how one might use it to make a claim about whether $B(p \ \& \ q) = (Bp \ \& \ Bq)$, by considering recent remarks by Chris Hill and Joshua Schechter:

¹⁸ See Heil 1989, pp. 571-583, for further (distinct) complaints about both the helpfulness and the intelligibility of the compartmentalization thesis.

Conjunction...can aggregate risk. A thinker may be justified in believing each of a set of propositions to a high degree but not in believing their conjunction to nearly as high a degree....

The general phenomenon of risk aggregation is well-known from discussions of the preface paradox. One version of this problem goes as follows: There are very many propositions that I count as knowing. Such propositions include simple claims of mathematics and logic; claims about myself, my environment, and my past experiences; and so on. Consider the conjunction of all of these claims. It seems that even were I cognitively able to competently deduce the conjunction from each of the individual claims that I know, I would not be in a position to know the conjunction. For I know that I sometimes – very rarely, perhaps – make mistakes. It should seem likely to me that at least one of the relevant propositions is false. It would be the height of arrogance to go on and infer the conjunction, knowing full well that there is a significant risk of falsity. As it happens, since the conjunction is a conjunction of claims that I know, if I were to draw the inference, I would infer a truth. But this would be a happy accident. The belief would not count as genuine knowledge. (2007, pp. 105-6)

As Hill and Schechter describe the lottery paradox, the paradox assumes, or at least relies on, the claim that there is a difference between believing some number of propositions and believing their conjunction. (The way the paradox is framed above also assumes, or again at least relies on, the idea that to believe a conjunction requires some kind of mental activity.¹⁹) For this reason, we could hardly use the paradox as worded to support the conclusion that $B(p \ \& \ q) \neq (Bp \ \& \ Bq)$; that would be question begging. But we can frame an argument that is based on the point Hill and Schechter explore in a way that one might think supports,

¹⁹ One wonders if Hill and Schechter are committed to the intelligibility of the following situation:

A subject believes that apples grow on trees and believes that oranges grow on trees, but decides to play it safe and refrain from believing that apples grow on trees and oranges grow on trees.

It is unclear how such a subject would answer the question: Is it the case that both apples and oranges grow on trees? If he says yes, then why be such a stickler about believing that apples grow on trees and oranges grow on trees? But what else could he say but yes? Would he say “I have no opinion on that question – I believe that each grows on trees, but I’m agnostic about whether they both do.” Perhaps this is an unfair way of putting it, but what is a fairer way of putting it? The trouble I have in understanding this scenario is not with imagining a subject with the willpower, or doxastic restraint, to refrain from believing that apples grow on trees and oranges grow on trees, given that he believes that apples grow on trees and believes that oranges grow on trees. I have trouble even imagining what such a belief state would be.

Moreover, would a cautious subject concerned to avoid conjunctive belief also refrain from believing what would be expressed by a person saying “the fact that apples grow on trees is not identical to the fact that oranges grow on trees”? I am inclined to think that she would not. Yet this belief seem to include as part of its content that apples grow on trees and oranges grow on trees.

rather than takes as a premise, that there is a distinction of the kind I am suspicious of. I will argue that even so framed, the line of argument is unpersuasive. The idea is this:

Suppose someone knows each of that p_1 , that p_2 , ... that p_n . This does not entail that he knows the conjunction that p_1 and p_2 ... p_n , for reasons of risk aggregation. Hence his being such that he believes each of that p_1 , that p_2 , ... that p_n is something other than his being such that he believes the conjunction – after all, he knows the former but not the latter.

The problem with this argument is that the conclusion does not follow. We can grant that knowing each of that p_1 ... that p_n does not entail knowing that p_1 and.... p_n without granting that this is because, or in any way shows that, the subject in question fails to believe that p_1 and.... p_n . For it may be that what risk aggregation affects is only whether one's believing that the world is a certain way meets whatever justificatory threshold is involved in knowledge.

By way of comparison, take a map of the world. It seems obvious that the likelihood that the world is as the map says it is is lower than the likelihood that a country picked at random is as the relevant part of the map says it is. But this does not show that the existence of the map as a whole, or of the map's having the content it has, is anything more than, or is not entailed by, each part of the map existing or having the content it does. If belief is relevantly similar, then it may be possible for a subject to believe that p , believe that q , and (thereby) believe that p and q , and to be such that although her believing that p and her believing that q counts as knowledge (of the fact that p and of the fact that q , respectively) her believing that p and q does not count as knowledge.

One might wonder if this conflicts with the basic worry that makes me suspicious of the claim that $B(p \ \& \ q) \neq (Bp \ \& \ Bq)$, which is that the fact that p and q is nothing more than the fact that p and the fact that q . How could one fail to know that p and q , if one knows that

p and one knows that q , if the conjunctive fact is nothing more than each of the facts? One would already know it, by virtue of knowing the facts it is nothing more than!

Why does the objection work in the case of belief but not in the case of knowledge? That is, why (as I am inclined to think) does the fact that the conjunctive fact that p and q is nothing more than the fact that p and the fact that q show that one believes that p and q whenever one believes both that p and that q , but not show, by the same measure, that a subject knows p and q whenever she both knows that p and knows that q ? The answer, I think, is something like this: When we say that someone knows a proposition, we *appraise* her believing of it, we assert that her believing of the proposition meets certain conditions. In general, a and b can entail c , in the sense that c is nothing more than a and b , without it being the case that every property attributed to both a and b is attributed to c . Consider the mereological sum of two atoms, each of which weighs one gram. The sum weighs two grams, not one. But this does not show that the atoms can exist without the sum existing (that might be the case, but it is not shown by the fact that the mass of the sum is greater than the mass of the parts). It seems to me that epistemic justification may be like this. The fact that my believing of p and my believing of q is each justified to extent e does not entail that my believing of p and q is justified to extent e . But that doesn't show that my believing of p and q is anything more than my believing of each of p and q .

Here is another analogy. Call an object K if it casts a light of less than 0.5 candlepower. Now imagine a screen composed of hundreds of pixels, each of which is capable of being illuminated to some degree (from 0 to 1 candlepower of light). The screen at t is such that every pixel is illuminated to brightness 0.2 candlepower. Each pixel would be K , but the screen as a whole would not be K , for it is casting far too much light. It would be K only when each pixel is casting only the tiniest amount of light, or one pixel is casting 0.4

and the rest casting none, etc., such that the aggregate amount of light cast by the screen is less than 0.5 candlepower. Again, if epistemic justification is like this (a point I have not even argued, let alone established) then the argument for $B(p \ \& \ q) \neq (Bp \ \& \ Bq)$ based on epistemic considerations would seem to fail.

There is a different way that one might put the problem that the lottery paradox, and similar puzzles, raise for the claim that $B(p \ \& \ q) \neq (Bp \ \& \ Bq)$, that calls for an additional response. Suppose one grants that a subject can be epistemically justified in believing each of that $p_1 \dots$ that p_n , but not epistemically justified in believing the conjunction that p_1 and... p_n . If one supplements this with a principle to the effect that one ought not to believe a proposition unless one is epistemically justified in doing so, then one might think there's a serious problem for the claim that $B(p \ \& \ q) \neq (Bp \ \& \ Bq)$. After all, how could a subject be normatively remiss in believing that $p_1 \dots p_n$ if she believes this just by virtue of believing each of that $p_1 \dots$ that p_n , given that her believing of each of those propositions is not normatively remiss? Or, to put it another way, how could it be the case that she ought to believe each of that $p_1 \dots$ that p_n , but ought not to believe that p_1 and... p_n , unless being in this last belief state is something other than, and moreover not entailed by, being in all the former belief states?

This is a good question. I am inclined to think what this shows is that the normative principle applies (if it applies at all) only to non-conjunctive beliefs. This might seem ad hoc – indeed, how could a solution be more tailored to the problem – but it is principled in fitting in with the general approach to conjunctive belief to which I am sympathetic. Believing a conjunction is not, on this picture, some additional, substantive state that a person may or may not enter into, once they are in the state of believing each conjunct. Thus, while we may be able to assess the epistemic merit of their believing of a conjunction

(that is, in regard to whether it meets whatever standard is required for knowledge), this does not give us reason to think that assessment is linked to a normative dictum concerning what cognitive behavior an agent ought to engage in.

Let me close this discussion by saying that these issues are puzzling indeed. It is worth noting that the epistemic argument considered, if successful against $B(p \ \& \ q) = (Bp \ \& \ Bq)$, would also appear to be successful against the possible worlds model of belief content. For Stalnaker, the main proponent of the view, holds that a belief state is closed under entailment, and that the believing of a “perfectly rational intelligence” (1984, p. 98) can be represented as a single belief state (that is, no compartments). Yet it seems that considerations of epistemic risk aggregation would nonetheless apply, such that the believing of some of the propositions in this set would count as epistemically justified, whereas the believing of others wouldn’t. If the epistemic considerations show, therefore, that $B(p \ \& \ q) \neq (Bp \ \& \ Bq)$, then they show the possible worlds model (at least as defended by Stalnaker) fails.

I do not take myself to have come anywhere close to establishing that to believe a conjunction just is to believe each conjunct. I am not convinced of that myself, of course. The point, instead, is to raise, and explore to the small degree possible here, some concerns that I have with the conjunctive proposal, which I said is perhaps the most natural way to understand the unity of belief. One such concern is, as I have just discussed at length, the fact that it is unclear how to square the common view that $(Bp \ \& \ Bq) \neq B(p \ \& \ q)$ with the equally common view that a conjunctive fact is nothing more than its conjuncts. There are two additional concerns I have with the conjunctive proposal, however, that are independent of this issue – two concerns that survive even if it turns out that there is a genuine distinction between $(Bp \ \& \ Bq)$ and $B(p \ \& \ q)$. In other words, and to emphasize, I think

there may be a problem with the conjunctive model even if there is a substantive or genuine distinction between $B(p \ \& \ q)$ and $(Bp \ \& \ Bq)$.

I can bring out these additional worries by considering a possible response to my skepticism concerning the commonly accepted view of conjunctive belief:

Suppose your worry about the alleged distinction between believing that p and that q and believing that p and q is well founded. This is no threat to either the conjunctive proposal or to the more general account you have offered. This is because if there is no difference between believing that p and that q and believing that p and q , as you have tentatively suggested, then you could still embrace both the general model and the conjunctive account. You would only have to accept that a subject always has an overall global representation (since she always believes her global conjunctive proposition).

This response might seem like a friendly amendment, but it reveals an important tension between the conjunctive model of global representation, the suggestion that $(Bp \ \& \ Bq) = B(p \ \& \ q)$, and the basic claim that motivates the project of this dissertation. Recall that the whole point, it seemed to me, is that having a global representation is not the same as merely having many beliefs. If this were not so, then there would be no point in asking whether, or under what conditions, believing subjects have an overall global representation, nor in asking for an explanation of the relation between a subject's many beliefs and her overall global representation. But if we accept that (i) there is no difference between believing some number of propositions and believing their conjunction, and that (ii) a subject has an overall global representation if she believes her global conjunctive proposition, then there seems to be little difference between believing many things and having an overall global representation. This way of putting it is a little too strong, for there might be, as I said earlier, ways of having an overall global representation (ways of meeting the general account) other than via conjunctive belief. So the distinction between having many beliefs and having an overall global representation would survive. But having many beliefs would be one way a

believer could meet the condition set by the general account, and so merely by having many beliefs a subject would have an overall global representation. And this conflicts with the basic motivating point that there is a problem to be understood concerning the nature of the global representations many philosophers take each of us to have and concerning the relationship between such a global representation and what we call a subject's individual beliefs.

The tension can be put starkly as follows:

- I do not see how there can be a distinction between believing that p and that q and believing that p and q
- If there is indeed no distinction, then (trivially) a subject always believes her global conjunctive proposition
- If a subject always believes her global conjunctive proposition, then a subject always has an overall global representation, on the assumptions that (i) the general model proposed above is correct, and (ii) believing a global conjunctive proposition meets that model
- Moreover, it would be merely by having many beliefs that a subject would have an overall global representation
- This conflicts with the claim that we require some kind of explanation or understanding of the unity of belief

The tension is obvious, and troubling. Three lines of resolution suggest themselves:

- 1. Deny that $(Bp \ \& \ Bq) = B(p \ \& \ q)$
- 2. Deny that there is a substantive difference between having many beliefs and having an overall global representation.
- 3. Deny that believing a global conjunctive proposition is all there is to a subject's having an overall global representation.

The first move may, in the end, be the right one to take: I have already discussed my reservations concerning it and will say no more about it here. The balance of the chapter will be devoted to arguing that there is good reason to think the proper line of resolution is 3.

The insufficiency of conjunction

Why 3? Note that there are two different ways to implement the suggestion in 3. One could agree that conjunction meets the general account proposed, and deny that the general account is correct. Or one could uphold the correctness of the general account and deny that conjunction is an instance of it. I will lay out my objections to conjunction, that is, my reasons for thinking option 3 is the better way to go, without saying whether the complaints target merely the conjunctive proposal or also the general account.²⁰

A first problem with the conjunctive model is that it is not clear that, even where a subject believes a conjunctive proposition, she thereby has one representation in the sense that seems to be involved when we speak of a subject having an overall representation. The problem is that a subject can believe a conjunctive proposition even in cases where, intuitively, it does not seem the subject's believing of each conjunct is tied together in the way required for the two beliefs to compose one representation. Second, there is reason to think that, whatever we make of conjunctive belief, the possibility of believing conjunctions (indeed, of believing in general) rests on something deeper, which is that one has sufficiently many 'diversity judgments' concerning the states of affairs one takes to obtain. 'Diversity judgment' is a technical notion, to be explained in more detail in a moment (and in yet more detail in the chapter to come). The basic idea behind this objection to the correctness of understanding global representation by appeal to conjunction is that one could not believe a

²⁰ It may seem obvious that if the conjunctive proposal fails, the general account fails, since the conjunctive proposal seems to be an obvious instance of that account. However, it is not clear to me that this is so; I think the issue turns on some difficult questions concerning the nature of conjunctive propositions. In any event, all I hope to do is express, and perhaps motivate readers to share to some degree, some further reservations I have with the conjunctive proposal. There is no need to decide at this point whether, if those reservations are good ones to have, the conjunctive proposal alone is inadequate or the general account is as well.

global conjunctive proposition at all if one did not believe a proposition of a different sort (i.e., a diversity proposition, in a sense to be explained), and that the believing of this proposition is a better candidate to explain the unity of belief.

Concerning the first of these points, that it is not clear that believing a conjunctive proposition is sufficient for a subject to have ‘one representation’. I don’t think the point can be demonstrated by describing a situation in which a subject believes her global conjunctive proposition but nonetheless fails (intuitively) to have an overall global representation. This is because, as I will argue in the chapter that follows, a subject necessarily has an overall global representation. Thus there are and can be no subjects who believe their global conjunctive proposition without having an overall global representation. However, I think the insufficiency of mere conjunction can be pressed by considering a subject’s beliefs concerning restricted domains, for there we can imagine cases where believing a conjunctive proposition and having an overall representation, in the sense that seems to be involved in the unity of belief, come apart.²¹

There are three kinds of cases to consider.

As an illustration of the first kind of case, suppose I believe that Henry Kissinger is a man and believe that Deep Throat is a man, but have no opinion about whether Kissinger’s being a man is the same as Deep Throat’s being a man (since I have no belief about whether Kissinger = Deep Throat). It seems to me that in this case my belief that Kissinger is a man and my belief that Deep Throat is a man do not fit together into one picture. Or, to put the

²¹ The examples appeal to intuition concerning whether a subject’s beliefs, or a subset of them, compose one representation, but the appeal to intuition is warranted here, I believe, since it is via intuition that we have our basic grip on the notion of an overall global representation. That is, the investigation into global representations rests on the fact that we have some intuitive or introspective sense that our beliefs form (or at least can form) such a representation. Hence if an imagined scenario seems to be one in which, intuitively, one does not have such a representation, then this is some evidence that the imagined scenario is not one in which the subject would have such a representation.

point more precisely, they may fit together into one picture (as constituents of some larger picture), but the two beliefs themselves, alone, don't fit together into one picture. And this is so even on the assumption that I believe the conjunction:

Henry Kissinger is a man and Deep Throat is a man.

rather than merely each conjunct on its own. To see this point, one has to imagine oneself into a situation in which one believes that Kissinger is a man and that Deep Throat is a man, but, because of some uncertainty concerning whether Kissinger just might be Deep Throat, one has no beliefs about whether Kissinger's being a man = Deep Throat's being a man. You might, for instance, in the course of wondering whether Kissinger = Deep Throat, notice that you believe both that Kissinger is a man and that Deep Throat is a man. And you might thereby wonder whether Kissinger's being a man, or that fact that Kissinger is a man, is identical to Deep Throat's being a man, or the fact that Deep Throat is a man. On this question, we are supposing, you have no opinion. It seems to me, intuitively, that in a situation like this, one's belief that Kissinger is a man and Deep Throat is a man (that is, one's conjunctive belief, which we can easily suppose one has, regardless of what we decide on the question of whether $(Bp \ \& \ Bq) = B(p \ \& \ q)$) does not have as its content a unified representation in the sense that we are trying to understand. One believes each thing (that Kissinger is a man and that Deep Throat is a man) but the two beliefs don't compose one picture. They are conjoined, we can grant, but for all that they are not a unity.

Analogously, imagine a case in which you are lost in a vast library. You have discovered, by carefully setting forth and then retracing your steps, two things. You have learned that if you head down the hall to the left you will eventually, after a huge number of complex twists and turns, go through a door that overlooks a large and busy atrium, full of people passing to and fro. And you have learned that if, having returned via the same route

to your original starting point, you head down a hallway to the right, in the opposite direction, you will eventually come, after another series of complex twists and turns, to a door that, when opened, overlooks a large and busy atrium, much the same – in fact, indiscernible so far as you can tell – from the atrium encountered via the other route. In this kind of case it seems quite reasonable to think that you would have no opinion concerning whether the two atriums are really the same atrium, arrived at via different routes, or whether they are in fact distinct atriums, though similar in appearance due to architectural uniformity and the uses to which they are put. In this situation, it seems to me, you would believe that ‘Atrium A is a busy place’ and that ‘Atrium B is a busy place’, but these beliefs would not fit together into one picture given your doubt about the identity relation between Atrium A and Atrium B and hence about the identity relation between the fact that Atrium A is a busy place and the fact that Atrium B is a busy place. And this is so, of course, even if you believe the conjunction ‘Atrium A is a busy place and Atrium B is a busy place’.

The second kind of case is a familiar one: consider Lois Lane and her representing of the entity that is Clark Kent/Superman. She believes that Clark Kent is a reporter, that Superman can fly, that Clark Kent lives in the city that she lives in, that Superman does too, and so on. However she also believes that, for instance, Clark Kent’s living in the city that she lives in is not identical to Superman’s living in the city that she lives in. (As she might put it, were she ever asked, “the fact that Clark Kent lives in the same city as me is not identical to the fact that Superman lives in the same city as me. They are two facts, not one.”) In this kind of case, it seems to me, Lois Lane’s beliefs about Clark Kent and Superman do fit together, along with her many other beliefs, to present her with one picture of the way things are, but they do not fit together to present her with one picture of that entity that is Superman/Clark Kent. She does not have one picture of that entity that is Superman/Clark

Kent even if we suppose she believes whatever conjunctive proposition that has as conjuncts each of the propositions that are the content of her individual beliefs about Clark Kent and Superman. (Note, though, that intuitively it seems that Lois Lane *does* have one picture of Clark Kent, and one picture of Superman.)

Finally, there is a third kind of case, one that involves neither the absence of a diversity judgment (as in the first case) nor a diversity judgment that mistakenly asserts a non-identity. Imagine that Smith is dating a set of identical twins without realizing that he is dating two people rather than one. Sometimes Smith hangs out with Robin, sometimes with her sister Sarah, and this continues for years without him being any the wiser. It seems to me that in this case Smith would have one overall representation (i.e., one picture) of an entity that does not in fact exist. But this is not because, of course, he has conjoined his many beliefs (whatever that might mean or amount to). After all, were he to be knowingly dating two people he would not have an overall representation, or one picture, of some entity that does not exist, not matter how much conjoining he did.

These three kinds of cases concern only the representing of objects (or perhaps any restricted domain), rather than all of one's beliefs in an unrestricted sense. But what they point to, I think, is that believing a conjunctive proposition is not sufficient for it to be the case that the content of what is believed is unified into one representation, in the sense we have in mind when we speak of an 'overall representation'. In the chapter that follows, I'll argue for a further claim, which can perhaps also be seen in these cases, which is that diversity judgments are what does the job of bringing unity. But that can wait: the point is only that the cases suggest that conjunctive belief is not sufficient for the having of one representation.

The final doubt I have with the conjunctive proposal is distinct from, but connected to, the doubt just discussed. It seems to me that whatever we make of conjunctive belief (that is, regardless of whether we end up deciding $B(p \ \& \ q) = (Bp \ \& \ Bq)$), the possibility of believing conjunctions rests on something deeper, which I call diversity judgments²². I will explain what I mean by a diversity judgment in more detail in the following chapter, but the idea, very roughly, is that a diversity judgment is a judgment of identity or non-identity between states of affairs asserted, by that judgment, to obtain. The best way to get a handle on the idea is simply to think of the judgment you would make were you to assert, for instance, any of :

The fact that Brown University is in Rhode Island is not identical to the fact that Queen Elizabeth is female.

The fact that no sheep has swum across the Atlantic is not identical to the fact that sometimes, in Spring, frost heaves buried rocks upward toward the surface of the soil.

The fact that Kissinger is a man is not identical to the fact that Deep Throat is a man.

That is, you believe that Brown University is in Rhode Island, and you believe that Queen Elizabeth is female. But you also believe something else, which you would express by saying “the fact that Brown University is in Rhode Island is not identical to the fact that Queen Elizabeth is female”. And so too with the other examples.

It does not appear to be metaphysically necessary that a subject believes a relation of non-identity obtains between any two metaphysically distinct states of affairs that she takes to obtain. That is why, for instance, it is possible to have no opinion on the question of whether Kissinger's being a man is identical to Deep Throat's being a man, even when one believes both that Kissinger is a man and that Deep Throat is. However, it also seems clear

²² I thank Jaegwon Kim for suggesting this expression.

that a subject could not believe her global conjunctive proposition – in fact, could not be a believer in any intelligible sense – unless she had many such beliefs concerning non-identity.

I cannot discuss this second point in as much detail as I would like, but very roughly, its intuitive appeal can be illustrated as follows. Is it possible for there to be a subject such that:

- (i) it is correct to attribute to the subject the belief that p_1 and p_2 and p_3 and p_4 and p_5 and p_6 , etc.
- and
- (ii) the subject has *no* beliefs of the kind she would express by saying “the fact that p_1 is not identical to the fact that p_2 ”.

It seems to me the answer is obviously no, for her beliefs would in some sense collapse. If she does not distinguish, for instance, the fact that p from anything else that she believes, then that the proposition p_1 is true is, for her, not something different from that the proposition p_2 is true, or from that the proposition p_3 is true, and so on.

The intuition might be stronger if we fill in the details rather than use schematic letters to stand for propositions. So let’s ask whether a subject – any subject – could believe that oranges are a kind of fruit, and that summers in Texas are hot, and that Pluto used to be classified a planet but isn’t any longer, and that cows have four stomachs, if she had no opinions of the kind she would express by saying:

the fact that oranges are a fruit is not identical to the fact that cows have four stomachs.

the fact that summers in Texas are hot is not identical to the fact that Pluto used to be classified as a planet but isn’t any longer.

the fact that summers in Texas are hot is not identical to the fact that cows have four stomachs.

etc.

It seems to me that if a subject did not believe the propositions in the second set, she could not be said to believe the propositions in the first set. If, for instance, I found out that some

person had no opinion on the question of whether the fact that summers in Texas are hot is identical to the fact that cows have four stomachs, I think I could safely conclude that she does not believe both things. In saying this, I am expressing sympathy with a general line of thought that is widely, albeit not universally, accepted. Stephen Stich aptly describes the core intuition:

Shortly before her death, Mrs. T had lost all memory about what assassination is. She had even forgotten what death itself is. She could, however, regularly respond to the question, “What happened to McKinley?” by saying, “McKinley was assassinated.” Did she, at that time, believe that McKinley was assassinated? For just about everyone to whom I have posed this question, the overwhelmingly clear intuitive answer is no. One simply cannot believe that McKinley was assassinated if one has no idea what an assassination is, nor any grasp of the difference between life and death. (1983, p. 56)

As Stich says, to believe (for instance) that McKinley was assassinated, one has to have an adequate understanding of what it is for the proposition ‘McKinley was assassinated’ to be true. And this, it seems to me, entails that if a subject believes that McKinley was assassinated, then the subject has many beliefs of the kind she would express by saying “the fact that McKinley was assassinated is not identical to the fact that p ”, where p is something else the subject believes. If the subject had no opinion on the question of whether (as she would put it) the fact that McKinley was assassinated is or isn’t identical to the fact that McKinley was a president, or the fact that assassination is a form of killing, or the fact that snow is white, etc., then the subject could not be said to believe that McKinley was assassinated.²³

It may also be possible to motivate the intuition the following way. Suppose you were charged with the task of picking a random person off the street and educating him or her about something unfamiliar to them. All you have to do is teach that person at least one

²³ The subject need not believe that the fact that McKinley was president is not identical to the fact that p , for every proposition she believes (otherwise Hesperus/Phosphorus cases would be impossible). But it seems clear that she must, in the main, distinguish it from the rest of what she believes.

thing she doesn't already know.²⁴ Would it be possible to teach her anything without teaching her a difference? That is, would you be able, even in principle, to bring it about that she has a belief she did not have before without teaching her a difference? Perhaps I should say more clearly that what I am asking the reader to consider is not whether one could teach her something without doing so via ostensive discussion of some difference. Rather, the question is whether it would be possible to teach her something without her coming thereby to acquire some new (or rather countless new) diversity judgments, that relate the content of her new belief to the bulk of the rest of what she believes. Suppose the new thing you manage to teach her is that you are a philosopher. It seems to me that if she did not grasp that that you are a philosopher is other than that apples are red, that New York is larger than Toronto, that one can tell the age of a tree by counting its rings, etc., (assuming she believes all these things), then it would be absurd to attribute the belief that you are a philosopher to her. Her coming to believe that you are a philosopher seems to involve her grasping that that you are a philosopher is other than innumerable other things that she believes.

Conclusion

The discussion in this chapter has been somewhat far-ranging, so let me close by reviewing the dialectic.

The chapter began by setting forth for discussion a two-part proposal concerning how we may be able to understand overall global representations. I offered a general formulation of the conditions a subject would have to meet if she is to have an overall global representation. This general account was complemented by a model that appealed to

²⁴ The wording here speaks of knowledge, but really what is at issue is merely belief. The idea is: all you have to do is get her to believe something she doesn't already believe.

conjunctive belief, which is perhaps the most natural way to understand the unity of belief. The idea is that by believing her global conjunctive proposition, a subject meets the condition offered by the general account. This first part of the chapter concluded with a brief discussion of how the two-part model on offer largely reduces the question of whether or under what conditions a human subject has an overall global representation to the question of whether or under what conditions a human subject believes her global conjunctive proposition; three different views on this latter question were canvassed.

The remainder of the chapter focused on articulating worries I have about the correctness of the conjunctive model. A first and important problem centers on the intelligibility of the alleged distinction between $(Bp \ \& \ Bq)$ and $B(p \ \& \ q)$. I argued that it is not clear that the distinction is sound, for reasons that have to do with the metaphysics of conjunctive facts, and briefly considered how one skeptical of the distinction may be able to respond to standard lines of thought that motivate accepting the distinction. I then noted that there is a tension between the schematic model/conjunctive proposal, the claim that $(Bp \ \& \ Bq) = B(p \ \& \ q)$, and the claim that the unity of belief is something that requires an explanation, in the sense that there is a substantive difference between merely having many beliefs and having one representation that these beliefs in some sense compose. In light of this tension, one could reject the two-part model, reject the claim that conjunctive belief is ‘automatic’, or reject the claim that the unity is something needing an explanation. The chapter closes by exploring two lines of arguments that point to independent problems with the conjunctive model, one having to do with conjunction being insufficient for unity as that is intuitively understood, the other saying that, whatever we make of conjunctive belief, the possibility of believing conjunctions, indeed of believing at all, seems to rest on something deeper, namely the having of ‘diversity judgments’.

In the next and final chapter I turn to exploring the potential for diversity judgments to explain the unity of belief.

UNITY AND DIVERSITY

If each representation were completely foreign to every other, standing apart in isolation, no such thing as knowledge would ever arise. For knowledge is a whole in which representations stand compared and connected

– Kant, *Critique of Pure Reason*, A98

There is something right about the general line of thought that motivates the conjunctive proposal, even if, as I have argued, there is reason to think the proposal inadequate. What is right is that the unity of belief cannot be understood simply by positing more beliefs (about, for instance, explanatory relations between propositions believed, or about inferential connections among beliefs) unless those additional beliefs in some appropriate sense subsume a subject's individual beliefs. For otherwise positing additional beliefs just adds more beliefs to the pile, as it were, leaving us with more items to unify.

This chapter focuses on exploring the idea that what 'binds together' a subject's individual beliefs is a diversity judgment, the rough idea of which was introduced in the previous chapter. I emphasize that what follows is a suggestion, an early pass at connecting the concept of a diversity judgment to that of the unity of belief, and at taking the former to be involved in explaining or grounding the latter. What are involved are some delicate problems in the philosophy of mind, philosophy of language, and metaphysics; would that my efforts here were not so clumsy.

The chapter proceeds as follows. First, I say more about what a diversity judgment is, and about why one might think that, at least in a restricted class of cases, a diversity

judgment can be understood as subsuming the content of other beliefs in a way that makes the subsumed beliefs ‘doxastic parts’ of the diversity judgment. The rough idea is that a subject’s having a belief of the kind she would express by saying, for instance, “the fact that p is not identical to the fact that q ” has the belief that p and the belief that q as ‘doxastic parts’. The goal of this first step is to propose a way of understanding how, or at least one of the ways in which, two or more beliefs can be ‘bound together’ into one representation. The idea is to replace talk of binding, fitting together, being tied together, constructing, building up, etc., with something more precisely understood. As Russell put it, the hope is to pass “from these obvious, vague, ambiguous things, that we feel quite sure of, to something precise, clear, definite, which by reflection and analysis we find is involved in the vague thing that we started with, and is, so to speak, the real truth of which that vague thing is a sort of shadow.” (1985, p. 179-180)

Having introduced the core idea, I turn to exploring how the concept of a diversity judgment might be used to ground the unity of belief, understood as a subject’s having of an overall global representation. That is, I look to broaden its application beyond the restricted class of cases considered in introducing the concept and illustrating how it is a way in which we can think of beliefs being ‘tied together’. I identify the central challenge in developing such an account to be the fact that subjects can suspend judgment on questions of the identity of states of affairs they believe to obtain. An example, used also in the previous chapter, is that one could believe that Deep Throat is a man, and believe that Kissinger is a man, but not have the diversity judgment one would express by saying “the fact that Deep Throat is a man is not identical to the fact that Kissinger is a man.” One therefore cannot have an overall global representation by virtue of necessarily having a belief state that affirms a relation of identity or diversity between every state of affairs one takes to obtain.

Having identified the central challenge, I argue it can, or at least can likely, be met: the chapter concludes with a simple and straightforward argument that I think gives us good reason to believe that a subject's beliefs form a unity in virtue of the subject having a diversity judgment that ranges over *all* of what we call her individual beliefs.

Diversity Judgments

At the center of my proposal is the concept of a *diversity judgment*. This is just a term of art for a belief state that affirms a relation of diversity or identity between two or more states of affairs the belief state affirms obtain. To get a better sense of what I am talking about, consider these two beliefs of mine:

B1: there are blackbirds

B2: the diameter of an electron is 10^{-18} metres

As a terminological point, B1 is the belief that there are blackbirds, not the content 'that there are blackbirds', understood as the sort of thing that could also be the content of a desire, or a doubt, or a command. I have both B1 and B2, but I also have a belief I would express in a sentence as:

B3: the fact that there are blackbirds is not identical to the fact that the diameter of an electron is 10^{-18} metres

Note that this is not the claim that I believe that the *beliefs* B1 and B2 are non-identical (although I do believe that as well). Rather, the idea is that in addition to believing that there are black birds and that the diameter of an electron is 10^{-18} metres, I take it to be the case that that it is the case that there are blackbirds is other than that it is the case that the diameter of an electron is 10^{-18} metres. This is a diversity judgment. Note also that B3 is distinct from the belief that the *propositions* 'there are blackbirds' and 'the diameter of an electron is 10^{-18} meters' are diverse. That belief simply asserts a relation of non-identity

between two propositions; in having it, a subject can remain neutral on the question of whether blackbirds exist, and on the question of whether the diameter of an electron is 10^{-18} metres. The belief B3, however, is not like this: it has as part of its content both that there are blackbirds and that the diameter of an electron is 10^{-18} metres.

We also have diversity judgments that assert a relation of *identity* between states of affairs affirmed by the judgment to obtain. For instance, I believe:

B4: William Felt is a man.

and

B5: Deep Throat is a man.

And I also believe something I would express by saying:

B6: The fact that William Felt is a man = the fact that Deep Throat is a man.

This is also, as I am using the term, a diversity judgment. It might seem odd to call this a diversity judgment, since it asserts a relation of identity rather than diversity. The main motivation for using 'diversity judgment' for beliefs that assert either identity or diversity between states of affairs affirmed by the belief state to obtain is the fact that *one* term was needed, and, moreover, the term 'identity judgment' is already widely used in philosophy to refer to a belief in the identity of objects (e.g., Hesperus = Phosphorus). But in any event, I am merely stipulating the meaning of a term, and so long as there is no risk of confusion, there should be little complaint about the choice of name.

What do diversity judgments have to do with the unity of belief? The relevance turns on the fact that B3 includes, within it, both B1 and B2. B3 is true only if both B1 and B2 are true. Moreover, any subject who has B3 necessarily has B1 and B2. In having B3, I believe both that there are blackbirds and that the diameter of an electron is 10^{-18} metres. As Davidson might put it, in having B3 a subject sees that there are blackbirds and that the

diameter of an electron is 10^{-18} metres ‘in a single glance’. In having B3, I pull what I believe in having B1 and B2 within the scope of a single judgment.

I propose that we think of B1 and B2 as ‘doxastic parts’ of B3. Recall this notion is:

A belief x with content Cx is a doxastic part of a belief y with content Cy iff (i) Cy entails Cx , (ii) necessarily, whosoever entertains Cy entertains Cx , and (iii) necessarily, whosoever has y has x .

It is clear, I think, that B1 and B2 meet the conditions set forth in this definition relative to the belief B3 (that is, that B1 and B2 are, by this definition, doxastic parts of B3).

What is the advantage of thinking of B1 and B2 as doxastic parts of B3? Moreover, what is the relationship between the concept ‘doxastic part’ and the more basic or familiar mereological notion of a part? There is much that can be said on both questions. Briefly, the advantage is that belief seems to have, as I have emphasized in Chapter One, something like a many/one structure, and an overall global representation seems to be a representation in some sense built up from a subject’s individual beliefs. The best, perhaps the only way, to make sense of this is by construing a subject’s beliefs as standing in a part/whole relation, or something very much like a part/whole relation. To be sure, it sounds odd to talk about beliefs as being parts of other beliefs; it is not clear what this means.¹ But beliefs have to stand in *something* like a mereological relation to one another if belief has, or even could have, the kind of unity it is alleged to have. The notion of a ‘doxastic part’ is conceived with this in mind. The three conditions given in the definition of ‘doxastic part’ are intended to

¹ This is perhaps because mereological relations are most familiar from discussion of material composition; this may lead us to think that if a is a part of b , a and b have to be little material chunks of some kind. But this of course is false: As Achille Varzi has emphasized, “it is worth stating explicitly that mereology assumes no ontological restriction on the field of ‘part’. The relata can be as different as material bodies, events, geometric entities, or geographical regions... as well as numbers, sets, types or properties.” (2003, S. 1) Moreover, there are ways of using mereological vocabulary in the context of belief that sound neither strange nor unfamiliar: For instance, one might say that part of what it is to believe that Smith is a bachelor is to believe that Smith is unmarried.

capture what it is for one belief to be part of another. When this is so, we can say the first belief is a doxastic part of the second.²

What is on offer so far is a first step – a way to understand the unity of belief in at least a restricted class of cases. That is, in place of talk of ‘binding’, ‘interlocking’, ‘fitting together’, ‘hanging together’, ‘being located in a network’, and the like, all of which are valuably suggestive but unsatisfying metaphors, I have offered an account on which two beliefs are doxastic parts of a third, where this notion is both (i) relatively clearly defined in familiar logical/metaphysical/epistemic vocabulary, and (ii) intuitively an instance of two beliefs being ‘unified’ into a third. This allows us, at least in the restricted domain I have discussed so far, to understand how belief can be both many and one.

At this point one might object and say “Big deal. Why think that diversity judgments are involved in a subject’s having a global representation of the world? Maybe diversity judgments can glue together two beliefs, but can they glue them all?” This is the key question. In reply, I think there is an intuitive connection between judgments of identity and diversity and the unity that seems to be involved in having a global representation. Moreover, I think it is possible to give an argument that shows, or at least makes it seem very likely, that one has an overall global representation that has as its content a diversity judgment.

Identity and Diversity

² There is more that should be said on this topic, at least at some point. An important thing to note is that belief contents may (and in fact seem to) have parts that are not doxastic parts. For instance, part of the content of the belief ‘Mary loves John’ is the content ‘Mary’. This is however not a *doxastic part*, since it is not a belief. (This is, at least, the standard view: one might wonder whether, if a descriptive theory of names is correct, and if, in general, the content of a concept is given by something like a theory, it may be ‘propositions all the way down’.)

In the previous chapter, I tried to give reason to think that mere conjunctive belief isn't sufficient for the contents of a subject's beliefs to be unified into one representation in the sense that seems intuitively to be involved in the having of a global representation. Someone who, for instance, has no belief she would express by saying "the fact that Kissinger is a man = the fact that Deep Throat is a man" cannot, and does not know how to, fit these two beliefs together; those beliefs do not, jointly, present her with one picture of the way things are. Her problem is that those facts are, for her, neither two facts nor one. Yet if those facts are, for her, neither two nor one, then in what sense could those beliefs be unified such that they present her with one picture of the way things are?

Broadening the account

The next step in the overall proposal is to see whether we have reason to think a diversity judgment account can be generalized beyond the restricted domain thus far considered. That is, we now have some clue about how it might be possible to understand what it is for beliefs to be unified, or to constitute a unity, in a restricted class of cases, but the discussion has been restricted to three specific beliefs I have. Is it possible to use this understanding of composition to make sense of the unity of belief writ large (or larger)?

In sketching how a diversity judgment seems to unify two (or more) beliefs, I referred to three rather simple beliefs. And I simply stipulated that, or ask you to take my word for the fact that, I believed the third. So the first order of business will be to assess what reason there is for thinking that one has more such beliefs.

To address this, we should note that there is nothing special about the beliefs I appealed to above. That is, they seem to be no different from the vast majority of my beliefs. Moreover, there is nothing special about me that makes it the case that these beliefs are

doxastic parts of some other belief. That is, on this particular issue I am no different from just about anyone else. In other words, it is at least typically the case that anyone who believes both that there are blackbirds and that the diameter of an electron is 10^{-18} metres believes something she would express by saying “the fact that there are blackbirds is not identical to the fact that the diameter of an electron is 10^{-18} metres”. It is not that I am smarter than they are and have figured out something that they have failed to.

One might object here that the relevant difference between me and someone who has beliefs B1 and B2 without having belief B3 is that I have taken the time to entertain the proposition. That is, one might object that Jones believes there are black birds and believes the diameter of an electron is 10^{-18} metres but does not believe what he would express were he to say “the fact that there are black birds is other than the fact that the diameter of an electron is 10^{-18} metres”, on the grounds that it has never occurred to him, or that he has never thought about it. One might even supplement this by telling a story in which Jones smacks his forehead one day, or otherwise manifests his realization of the obvious, thereby proving he had hitherto been ignorant of it. I take this to be an implausible objection, one that misunderstands the nature of belief. There are all kinds of things I believe that I have never thought about. For instance, I believe, as I did yesterday, that George Bush has never been at the center of the sun. I had never entertained that proposition before, but I nonetheless took the world to be such that George Bush had never been at the centre of the sun. I did not leave that open.³

Not only is there nothing special about me, there is nothing special about the beliefs considered so far. Most of my beliefs seem to be relevantly similar. Consider my believing:

Some coffee cups are made of paper.

³ More could be said on this issue, to be sure. If a reader is concerned at this junction, they can treat the points I'm making as conditional on a proof that one need not have considered a proposition in order to believe it.

It is often possible to determine the age of trees from growth rings.
 Chile is a mountainous country.
 Shoes were invented a long time ago.
 One day contact lenses will be obsolete.
 More people live in New York than in Toronto.
 The function of the heart is to circulate blood.
 etc.

The idea is that if I sit and generate sentences that express my beliefs, a great many of these beliefs are such that, as it happens, I have the relevant diversity judgment. That is, my believing that some coffee cups are made of paper and that Chile is a mountainous country involves (or is accompanied by, to leave things as neutral as possible) my believing what I would express by saying “the fact that some coffee cups are made of paper is not identical to the fact that Chile is a mountainous country”. So it is for the rest of the beliefs listed and countless more. Moreover, not only do I believe something I would express by saying:

the fact that Chile is mountainous is not identical to the fact that New York is larger than Toronto

I believe something I would express by saying:

the fact that Chile is mountainous is not identical to the fact that New York is larger than Toronto, and neither of these facts is identical to the fact that it is often possible to determine the age of trees by counting their rings.

And, similarly, I believe something I would express by saying:

the fact that Chile is mountainous is not identical to the fact that New York is larger than Toronto, and neither of these facts is identical to the fact that it is often possible to determine the age of trees by counting their rings, and none of these facts is identical to the fact that it is the function of the heart to circulate blood.

And so on.

So far, the claim is a restricted one: that a great many of my beliefs, and those of other people, happen to be such that they are doxastic parts of diversity judgments. This is somewhat hard to establish conclusively, for it would seem to involve checking belief pairs (or triples, quadruples, etc.), calling up two candidate beliefs and asking myself whether I

take (or prior to my calling up took) it to be the case that the fact that the first is true (to put it roughly) is not identical to the fact that the second is true. One could obviously not complete the testing – it just is not possible to call all of one's beliefs to mind, two at a time, for the purpose of such testing. But although one could not complete the testing, one can call to mind a great number of examples and test those. If, as I believe, the great majority of those will be doxastic parts of some diversity judgment, this is some evidence that the great majority of one's beliefs in general are doxastic parts of a relevant diversity judgment (some diversity judgment or other). This does not seem to be so merely for reasons of numerical induction: it is so not merely because if one calls to mind 50 belief pairs and finds them all doxastic parts of relevant diversity judgments, one infers that the most of the rest are likely to be too. Rather, what happens is that in calling to mind belief pairs and considering whether one has a corresponding diversity judgment, one comes to understand why it is that one overwhelmingly tends to have such diversity judgments. That is, one comes to suspect, I think, a stronger point, which is that it not merely *happens* to be the case that many of one's beliefs are doxastic parts of a diversity judgment, but that they must be. That is, it simply couldn't be the case that, in the main, one failed to have diversity judgments. It is hard to articulate clearly why this is so. But perhaps I can say something to illustrate that it is so.

By way of comparison, it seems a competent speaker of English could fail to have an opinion on whether two English words have the same meaning; for instance, he might have no opinion on the question of whether the term 'tidal marsh' and 'salt marsh' are synonymous. But no one could be in such a state in regard to every English word, and still count as a speaker of English. (Imagine a dictionary that, under every lexical entry, simply said: 'may or may not be synonymous with e1, e2, e3, e4, etc.', where each of these are other

entries in the dictionary.)⁴ The problem in these cases is that it isn't possible to make sense of such a would-be speaker of English understanding anything, or of the 'language' in the dictionary as saying anything or of being used to say anything.

Now suppose someone failed to have any diversity judgments – they believed lots of things, but had no opinion on whether the things believed were identical or diverse. Such a person's beliefs would fail to group things into classes, make distinctions, ascribe certain properties to some things and certain other properties to other things. In short, they would fail to be beliefs.

The central challenge to a diversity judgment account of unity

I have articulated a vague modal claim – many of one's beliefs must be doxastic parts of relevant diversity judgments. I won't refine or distinguish this claim from others, for what matter is only that we see that a certain kind of modal claim cannot ground the unity of belief. An option that *will not* work: It cannot be necessary that a subject has a diversity judgment that asserts a relation of diversity or identity between each thing that she believes, along the lines of what she would express by saying:

The fact that p1 is not identical to the fact that p2, neither of these is identical to the fact that p3, none of these are identical to the fact that p4, and so on.

This is not because such a proposition would be too 'complicated' to entertain; as I argued in the previous chapter, I do not think there is good reason to think a subject has to be able

⁴ Locke makes a point similar to this in his *Essay Concerning Human Understanding*, when he identifies as an operation of the mind the capacity to discern, in the sense of distinguish, ideas, and argues that there is "no knowledge without it [discernment]. There, talking of *ideas*, he says:

Another faculty we may take notice of in our minds is that of *discerning* and distinguishing between the several *ideas* it has. It is not enough to have a confused perception of something in general. Unless the mind had a distinct perception of different objects and their qualities, it would be capable of very little knowledge, though the bodies that affect us were as busy about us as they are now and the mind was continually employed in thinking." (Locke, *Essay*, Book II, Chapter XI, para. 1)

Locke is of course talking about ideas, and his particular empiricist understanding of them, but the general point motivating what he is saying is the point, I think, that motivates what I am saying. If a subject failed to discern, or distinguish, any given idea from any other ideas, then in some sense there would be no difference between her having all of the ideas and her having none of them.

to entertain a proposition in order to believe it. Rather, it cannot be necessary that a subject believe such a proposition because, as I have said before, it seems possible for subjects to be such that they hold open the identity of states of affairs they take to obtain (or, to put it more precisely, to be such that, where they to report on their relevant doxastic state, they would say: “I have no opinion on whether the fact that p = the fact that q ”). This is what one does who believes that Kissinger is a man and that Deep Throat is a man but has no belief concerning whether Kissinger's being a man = Deep Throat's being a man. The challenge, then, is to see whether diversity judgments can deliver the unity of belief even though such suspended judgments are possible. I think the answer is yes.

The first thing to note is that this kind of case is, it would seem, rather rare. Suspending a diversity judgment is, I think, derivative of – or at least requires – suspending a judgment on the identity of objects, properties, events, propositions, and so on⁵. For instance, it seems natural to say that people used to believe that the Evening Star was not identical to the Morning Star and therefore something that they would have expressed by saying “the fact that the Evening Star is a heavenly body is not identical to the fact that the Morning Star is a heavenly body”. Cases in which a diversity judgment is suspended arise only in special circumstances in which evidence is arranged so that no belief either way about the identity of two entities is generated. Here it is interesting to note that standard examples that appeal to ignorance of identities involve circumstances in which a person seems to have a *false* belief about an identity rather than *no belief* about an identity. In other words, what is standardly called “ignorance of identities” is ignorance by virtue of having a false belief

⁵ A person who believes that Superman = Clark Kent, or that the proposition ‘snow is white’ is not identical to the proposition ‘sheep are mammals’, or who has no opinion on the question of whether pain = c-fibre firings, is judging (or in the last case, refraining from judging) whether two things are identical, but the judgment does not include within it that the objects, events, states of affairs, etc., obtain or occur.

rather than by virtue of having no belief.⁶ In any event, it seems right to say that subjects do not seem to often hold open identities concerning the entities that feature as objects of their thought and that, similarly, or consequently, they do not seem to often hold open identities of the states of affairs they take to obtain.

If this is right, it suggests that although cases like the one involving Kissinger and Deep Throat are possible, they are atypical. Their mere possibility would show that the unity of belief could not be grounded in any necessity to having a diversity judgment of the form sketched, asserting a relation of identity or diversity between every state of affairs the subject takes to obtain. But this does not show that the unity of belief cannot be grounded in a diversity judgment, and perhaps even a necessary type of diversity judgment at that.

The Argument

I think it is possible to give a rather simple – surprising simple – argument for the unity of belief that rests on the notion of a diversity judgment. In an earlier, much more complicated, and four-fold-long version of this chapter, I worked through a series of distinct arguments, considering a variety of situations which a subject could be in vis-à-vis the having of diversity judgments (situations in which a subject contingently suspends none, suspends some, suspends many). The advantage of this method of presentation was that it accustomed the reader, through longer and more detailed exposure, to the picture of belief being developed. It also gave a clearer sense of the compositional structure of a subject's overall global representation. But I now think that there is a simple, more straightforward argument and that the best way to test it is to send it forth on its own, rather than flanked by legions of arguments assailing a common target.

⁶ Lois Lane is ignorant of an identity, for instance, not because she fails to have an opinion on the question of whether Clark Kent = Superman, but because she believes Clark Kent is not identical to Superman. She is not merely ignorant; she is mistaken.

Consider what has to be true for a subject to fail to have an overall global representation: For a subject to fail to have one representation that has as its content everything that she believes, she has to have (at least) two beliefs the content of which does not fall within the scope of a single judgment.⁷ In other words, she has to believe something p where p is not part of the content of any other belief she has. Why is this? Well, if p is part of the content of some other belief, then that second belief subsumes the belief that p . Suppose the content of that second belief is q , where p is a proper part of q . Then either (i) q is *all* the subject believes, in which case she has one belief that has as its content everything that she believes, or (ii) she believes q and something else too, some proposition that is not a part of q . If (i) is the case, then the subject's believing is a unity. If (ii) is the case, then either (iii) she has some other belief r such that it has q as a proper part, or (iv) she does not. If (iii) is true, then we can repeat the process, asking whether she has any belief that has r as a proper part of its content, until we reach a point at which we are forced to choose between something like (i) and something like (iv). That is, at the end of the day, if (say) the subject has the belief that p , then either that is all she believes, in which case we have the required unity, or she has some other belief the content of which does not have p as a proper part (and is not itself a proper part of p). The former disjunct involves unity, so the threat is the second disjunct.

The point, for all its simplicity, can be somewhat hard to see, so it is worth comparing the physical analog. Suppose I tell you that I have inside a box some number (greater than zero) of physical entities, but I don't tell you whether the entities stand in part/whole relations such that there is in the box exactly one whole of which everything else in

⁷ This is not quite right, for if a subject had no beliefs at all then she would not have a single representation either, even though she wouldn't fail to have a single representation by virtue of having two or more beliefs that fail to fall within the scope of a single judgment. So let us stipulate the obvious: that we are considering only those subjects who do not fail to count as believers.

the box is a proper part. What has to be the case for there to fail to be one whole of which everything else is a proper part? Well, the box has to contain two (or more) entities that are not themselves proper parts of anything. If I tell you that there are two (or more entities) that are not themselves proper parts of anything, then you can infer that the box does not contain exactly one whole of which everything else in the box is a proper part. If, conversely, I tell you that it is not the case that there are two (or more) entities that are not themselves proper parts of anything, then you can infer that there is exactly one whole of which everything else in the box is a proper part.

For a subject's believing to fail to be unified, then, a subject has to have a belief the content of which is (i) not all that she believes, and (ii) not part of the content of any other belief she has. This is what it is for a person's beliefs to fail to be unified; it is, that is, what it is for a person to be such that what she believes does not fall within a 'single glance'.

Could there be such a subject? I am inclined to think that there could not be. She would have to believe something, say that p , believe other things (other propositions that are not proper parts of p , say q , r , s , etc.), and have no opinion whatsoever of the kind she would express by saying "the fact that p is not identical to that fact that q ", or "the fact that p is not identical to the fact that r ", etc. That is, p 's being the case would have to be, for her, neither the same as nor different from q 's being the case, r 's being the case, s 's being the case, and so on. It is very hard to see that anyone could meet this description.

A conditional conclusion

Let me summarize the argument, put the point as plainly as possible, and clarify the scope and limits of its conclusion. Recall that for a subject to fail to have an overall global representation – for a believing subject to fail to have one belief that has as its content

everything she believes – the subject has to have two (or more) beliefs the content of which is not the proper part of any other belief she has. For this to be so, the subject has to have some belief the content of which is a mere proper part of what she believes, and yet have no belief that has that content as a proper part. And this, it seems to me, requires that the subject take something to be the case (say, that p) where this is not all that she takes to be the case and yet she does not in any way compare or contrast this with anything else that she takes to be the case. That is, so long as a subject does not have a belief the content of which she fails to distinguish from the content of any other belief she has, she will have one judgment that includes within its content everything that she believes. This is very far from standard Kissinger/Deep Throat cases; there the subject has no opinion on whether Kissinger's being Φ is identical to Deep Throat's being Φ , but has all kinds of beliefs, it seems right to think, contrasting Kissinger's being Φ with various things and all kinds of beliefs contrasting Deep Throat's being Φ with various things. For instance, although she does not have the belief she would express by saying “the fact that Kissinger is a man is identical to the fact that Deep Throat is a man”, she does have the beliefs she would express by saying “the fact that Deep Throat is a man is not identical to the fact that lettuce is a plant”, “the fact that Deep Throat is a man is not identical to the fact that Nixon is president”, “the fact that Kissinger is a man is not identical to the fact that most birds can fly”, and so on. What it requires instead is a very radical bifurcation of a subject's doxastic commitments, such that the subject simply leaves open whether ‘the fact that p ’, where p is what is true according to one body of belief, is other than ‘the fact that q ’, where q is what is true according to the other body of belief.

I am not sure what to make of this as a possibility: it seems impossible, but I have, for the moment, no argument for that claim. For this reason, we should accept that the

account considered so far only gives reason to accept the unity of belief as a contingent claim: If the kind of bifurcation just described is possible, then, contingently, our believing is one (on the assumption that such bifurcation is very far from our doxastic experience). If such bifurcation is impossible, then it is necessary that our believing is one.⁸

The argument, in a sense, inverts the presumed dialectical structure. Normally one would think that the burden of proof is on someone who wants to claim a person has some giant belief that subsumes every other belief she has. How are you going to argue for *that*? But now the situation is like this: I have identified a condition that must hold if belief is to fail to be unified. A subject has to have at least one belief such that (i) its content is not all that she believes, and (ii) she has no opinion concerning whether ‘the fact that *p*’, where *p* is the content of that belief, is identical to the facts that *q*, that *r*, or that *s*, where these are each of the other things that she believes. There is a serious question of the possibility of a subject having such a belief.

It may be helpful to put the point dialectically. Suppose I claim that my interlocutor has an overall global representation. He denies this, and asks for a proof. I then point out that if he fails to have an overall global representation, then he must have some belief that is not the proper part of any other belief he has (and, of course, is not such that it has as its content everything that he believes). I then ask him to specify some belief of his that meets this condition. Suppose he nominates the belief that *p*. I then ask him whether he has any belief of the form “the fact that *p* is not identical to the fact that *r*”, where *r* is something else that he believes. If he grants that he does have such a belief, then his belief that *p* is, contra

⁸ Objection: Such diversity judgments would quickly become too complex to entertain, and hence to believe. How would a person ever have formed all these beliefs? It just seems crazy to imagine a person has all these beliefs. Reply: There is a substantive question about whether belief states should be understood as states that, necessarily, are available to introspection, result from some kind of deliberative activity, or that require a certain amount of conscious involvement on the part of the subject. I have already said that I don’t find this way of thinking about belief promising, and cannot revisit the issue here.

to what he had claimed, a proper part of some other belief he has. But if he professes to have no diversity judgment of the form given immediately above, then his believing p seems unintelligible. He wants me to grant that he believes that p , even though he does not distinguish p 's being the case from anything else that he believes.

Conclusion

I will close by highlighting just a few of the many issues that require more extensive investigation, if we are to be clear about the potential for diversity judgments to ground the unity of belief.

First, what exactly is the content of diversity judgments? The belief I would express by saying “the fact that ‘Toronto is a city is not identical to the fact that sheep are mammals” is clearly different from the belief I would express by saying “the proposition ‘Toronto is a city’ is not identical to the proposition ‘sheep are mammals’”. The difference is that the first belief, but not the second, asserts that Toronto is a city and that sheep are mammals. That is, it asserts a relation of diversity between states of affairs that it asserts obtain. The difference can perhaps be more easily seen if we think of propositions whose truth-value we have no opinion on. For instance, compare the belief someone would express by saying “the proposition ‘Mike has three sons’ is not identical to the proposition ‘Mike has three daughters’” with the belief someone would express by saying “the fact that Mike has three sons is not identical to the fact that Mike has three daughters”. Anyone who is a competent speaker of English is in a position to have the first belief, assuming they can use the name ‘Mike’ in a way that picks him out. But only someone who believes both that Mike has three sons and that he has three daughters is in a position to have the second belief. A diversity judgment, therefore, is not a judgment that merely asserts a relation of diversity or identity

between propositions; it includes as part of its content that the states of affairs between which identity or diversity is asserted do indeed obtain. So much is reasonably clear; but this does not mean that the content is well understood.

Readers may have noticed that for the most part I have specified the content of a diversity judgment indirectly, by appeal to a belief report given by the believer in question.⁹ For instance, I tried not to say: the subject believes that the fact that there are blackbirds is not identical to the fact that the diameter of an electron is 10^{-18} metres. Rather, I said: the subject believes what she would express by saying “the fact that Kissinger is a man is not identical to the fact that black birds fly”. Had I used the former locution, then *I* would have been asserting that there are two facts – that there are blackbirds and that the diameter of an electron is 10^{-18} metres – and that the subject in question believes these two facts are not identical. And that, of course, cannot be the right way to specify the content of the subject's diversity judgment – it isn't that there are two facts the subject believes to be non-identical, it is that the subject believes these are two facts and that they are non-identical, or something close to that.

The problem can perhaps be more easily seen if we think of a subject who has beliefs we think are false. For instance, think of a subject who has a belief she would report by saying “the fact that François Mitterrand is (now, in 2008) president of France is not identical to the fact that George Bush is (now, in 2008) the president of the United States”. Few subjects (I hope) have such a belief, since Mitterrand has been dead since 1996. But it certainly seems possible that someone could think Mitterrand is the president of France, in which case, on the assumption he also believes George Bush is the president of the United States, it is almost certain that he would have the diversity judgment mentioned above. I

⁹ I say 'for the most part', because in a few instances doing so would have been so cumbersome that it would have impaired rather than aided understanding.

could not specify the content of this diversity judgment of his by saying he believes the fact that Mitterand is the president of France is not identical to the fact that Bush is the president of the US. After all, Mitterand is not the president of France, so that sentence, if spoken by me, would seem to suffer a reference failure. Moreover, in saying that sentence I would express (or imply) something I don't want to express (or imply), for in saying it I would assert that Mitterand is the president of France. The diversity judgment is, nonetheless, perfectly intelligible to us: we just have to specify its content by saying the subject believes something he would express by saying “the fact that Mitterand is the president of France is not identical to the fact that Bush is the president of the United States”. So we do not have, at this point, a third-person vocabulary to directly express the content of a subject's diversity judgment.¹⁰

This raises special difficulties since there are people (e.g., philosophers) who disavow the existence of facts, and people (such as young children) to whom one might be reluctant to attribute the concept of a fact.¹¹ Do these people have diversity judgments? If so, what is the content of their diversity judgments, given that we cannot happily specify the content of their diversity judgments by pointing to propositions they would express by saying something like “the fact that p is not identical to the fact that q ”. There are two distinct problems here. The first concerns whether the account of an overall representation of the world I propose requires that a subject have the concept 'fact'. The second concerns how to specify the content of diversity judgments in those cases in which the subject would be unwilling or unable to employ 'fact' terminology (regardless of whether the subject has the

¹⁰One possibility is that a diversity judgment should be analysed as a special kind of conjunctive belief, perhaps roughly of the form “Toronto is a city and sheep are mammals and the propositions ‘Toronto is a city’ and ‘sheep are mammals’ are diverse”? It is not clear to me, however, that this analysis captures what I believe when I have a diversity judgment. But then what is the correct analysis?

¹¹I am grateful to Chris Hill for pressing me on this point and helping me see its significance.

concept 'fact'). I take both problems to be part of the more general problem of understanding more clearly the content of a diversity judgment.

A second, distinct problem concerns whether the necessity of diversity judgments iterates beyond first-order beliefs. I have said that, for instance, a person couldn't believe that Kissinger is a man without grasping that this fact (to put it roughly, of course) is different from the fact that p , that q , that r , and so on, where these are other things that she believes. But can she have the belief she would express by saying "the fact that Kissinger is a man is not identical to the fact that p ", without grasping that *this* is not identical to (speaking roughly, again) the fact that r , that s , that t , and so on? If the answer is no, we seem to have hit on a regress, and the number of beliefs the subject is required to have explodes to infinity. If the answer is yes, however, then one might wonder why it is possible to have these second-order beliefs without having corresponding diversity judgments, if indeed it isn't possible to have first-order beliefs without having corresponding diversity judgments. I am inclined to think the subject is not required to have these second-order diversity judgments, and that this can be defended in a principled way. (The idea, in a nutshell, is that diversity judgments differ in kind from first-order beliefs. They arrange, or fall out from the arranging of, what is asserted by first-order beliefs.) This would require, of course, considerable elaboration and defense; the point is that the general challenge concerning whether the requirement iterates needs to be addressed.

A third problem that needs closer examination concerns the relationship between diversity judgments, as defined in this dissertation, and identity judgments, understood along the lines of 'the Morning Star = the Evening Star'. That there are connections is obvious; what exactly those connections are is much less so.

Fourth, can I say more to establish that it is just isn't possible for a subject to have a (first-order) belief the content of which she fails to distinguish from the content of any other (first-order) belief? Conversely, if this is possible, then why does it seem so problematic to imagine the widespread, perhaps even systematically pervasive, suspension of diversity judgments? It should also be possible to say more clearly why one could not pervasively suspend judgments concerning the identity of the objects and properties for which we have concepts.

Fifth, the discussion in this chapter has largely proceeded independently of the extensive discussion of identity judgments that has already occurred in the analytic tradition. I have made no attempt to reconcile the remarks made here with, or even position them within, this extensive body of literature. I think this approach was justified, given the nature of the project. But for the diversity judgment account of the unity of belief to be properly developed, this would need to be done.

Finally, it would be valuable to investigate more closely various other kinds of mental or propositional unity. I have pointed out, in earlier chapters, that the unity of belief appears to parallel in certain ways the unity of consciousness and the binding problem. But it would be good to know in exactly which ways. Similarly, the question of what it is to have a global representation, perhaps of the world but certainly of such entities as Toronto, my mother, the truths of metaphysics, the mind-body problem, etc., seems similar to the kind of unity that is involved in a concept. (My overall representation of my mother seems to be built out of all those things I believe about her; my concept 'mother' seems to depend somehow on what I believe about mothers.) I have only lately come to appreciate this latter similarity, and hence have done, I regret, no work to explore it within this dissertation.

When you pass from the vague to the precise by the method of analysis and reflection that I am speaking of, you always run a certain risk of error. If I start with the statement that there are so and so many people in this room, and then set to work to make that statement precise, I shall run a great many risks and it will be extremely likely that any precise statement I make will be something not true at all. So you cannot very easily or simply get from these vague undeniable things to precise things which are going to retain the undeniability of the starting-point.

– Bertrand Russell, *The Philosophy of Logical Atomism*

BIBLIOGRAPHY

- Achinstein, P. (2001). *The Book of Evidence*, Oxford: Oxford University Press.
- Alechina, N., & Logan, B. (2001). Logical omniscience and the cost of deliberation. *Proceedings of the 8th International Conference on Logic for Programming, Artificial Intelligence and Reasoning*, Springer Lecture Notes in Artificial Intelligence, Vol. 2250, pages 100—109.
- Anderson, M. and Rosenberg, G. (unpublished). Content and action: the guidance theory of representation.
- Audi, R. (1982). Believing and affirming, *Mind*, 91:361, pp. 115-120.
- Audi, R. (1998). *Epistemology: A Contemporary Introduction*, London: Routledge.
- Axel, R. (2004). Nobel Lecture: *Scents and Sensibility: A Molecular Logic of Olfactory Perception*. Dec. 8, 2004, at Sal Adam, Berzeliuslaboratoriet, Karolinska Institutet.
- Baker, L. R. (1987). *Saving Belief: A Critique of Physicalism*. Princeton, NJ: Princeton University Press.
- Baldner, K. (1996). Subjectivity and the unity of the world. *The Philosophical Quarterly*, 46: 184, 333-346.
- Bartelborth, T. (1999). Coherence and explanations. *Erkenntnis*, 50: 2-3, 209-224.
- Bickerton, D. (1990). *Language and Species*. Chicago: University of Chicago Press.
- Bilgrami, A. (1998). Why holism is harmless and necessary,” in Tomberlin, (ed.), *Philosophical Perspectives 12: Language, Mind and Ontology*. Boston: Blackwell.
- Block, N. (1986). Advertisement for a semantics of psychology, *Midwest Studies in Philosophy*, 615-78.
- Block, N. (1995). An argument for holism, in *Proceedings of the Aristotelian Society, New Series*, Vol. XCIV, 151-169.
- BonJour, L. (1985). *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- BonJour, L. (unpublished). Analytic philosophy and the nature of thought.

- Braddon-Mitchell, D. & Jackson, F. (1996). *The Philosophy of Mind and Cognition*. Oxford: Blackwell.
- Brandom, R. (1994). *Making it Explicit*, Cambridge, MA: Harvard University Press.
- Bratman, M. (1992). Practical reasoning and acceptance in a context. *Mind*, 101:401, 1-15.
- Brazier, F.M.T., Dunin Keplicz, B.M., Jennings, N.R. and Treur, J. (1997). DESIRE: Modelling multi-agent systems in a compositional formal framework, *International Journal of Cooperative Information Systems*, special issue on Formal Methods in Cooperative Information Systems: Multiagent Systems, edited by M. Huhns and M. Singh, 6:1, pp. 67-94.
- Brook, A. (1994). *Kant and the Mind*. Cambridge, UK: Cambridge University Press.
- Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence* 47, 139-159.
- Bykvist, K. and Hattiangadi, A. (2007). Does thought imply ought? *Analysis*, 67:296 , 277–285.
- Chisholm, R. (1986). *Brentano and Intrinsic Value*, New York: Cambridge University Press.
- Chopra, S., Ghose, A., & Meyer, T. (2003). Non-prioritized ranked belief change. *Journal of Philosophical Logic* 32, 417-443.
- Churchland, P.M. & Churchland, P.S. (1983). Stalking the wild epistemic engine?. *Noûs* 17, pp. 5-18.
- Churchland, P., & Sejnowski, T. (1990). Neural representation and neural computation. *Philosophical Perspectives: Vol. 4, Action Theory and Philosophy of Mind*, pp. 343-382.
- Clarke, M. (2004). *Reconstructing Reason and Representation*. Cambridge, MA: MIT Press.
- Cohen, S. (2002). Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research*, 65:2, 309-328.
- Currie, G. (1986). Untitled review of *Inquiry*. *The Philosophical Quarterly*, 36:145, pp. 569-571
- Davidson, D. (1978). What metaphors mean. *Critical Inquiry*, 5:1, pp. 31-47.
- Davidson, D. (1979). The method of truth in metaphysics, in P. French *et al.* (eds.) *Contemporary Perspectives in the Philosophy of Language*, Minneapolis: University of Minnesota Press.
- Davidson, D. (1982). Rational animals. *Dialectica*, 36:4, 317-332.

- Davidson, D. (1986). Empirical content. In E. Lepore (ed.), *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson* (pp. 320-332). Oxford: Basil Blackwell.
- Davidson, D. (1991). Three varieties of knowledge. In A. P. Griffiths (ed.), *A. J. Ayer: Memorial Essays. Royal Institute of Philosophy Supplement 30* (pp. 153-166). Cambridge, UK: Cambridge University Press.
- Davidson, D. (1999). Reply to Nagel. In L. E. Hahn (ed.), *The Philosophy of Donald Davidson* (pp. 207-8). La Salle, IL: Open Court.
- Davidson, D. (2001a). The method of truth in metaphysics. In *Inquiries into Truth and Interpretation* (pp. 199-214). Oxford: Oxford University Press.
- Davidson, D. (2001b). Thought and talk. In *Inquiries into Truth and Interpretation* (pp. 155-171). Oxford: Oxford University Press.
- Davidson, D. (2001c). Truth and meaning. In *Inquiries into Truth and Interpretation* (pp. 17-36). Oxford: Oxford University Press.
- Davidson, D. (2001d). Radical interpretation. In *Inquiries into Truth and Interpretation* (pp. 125-139). Oxford: Oxford University Press.
- Davidson, D. (2001e). Introduction. In *Inquiries into Truth and Interpretation* (pp. xv-xxiii). Oxford: Oxford University Press.
- Davidson, D. (2001f). A coherence theory of truth and knowledge. In *Subjective, Intersubjective, Objective* (pp. 137-153). Oxford: Oxford University Press.
- Davidson, D. (2001g). Irreducibility of the concept of the self. In *Subjective, Intersubjective, Objective*, (pp. 85-93), Oxford: Oxford University Press.
- Davidson, D. (2004a). The problem of objectivity. In *Problems of Rationality* (pp. 3-18). Oxford: Oxford University Press.
- Davidson, D. (2004b). Expressing evaluations. In *Problems of Rationality* (pp. 19-38). Oxford: Oxford University Press.
- Davidson, D. (2004c). What thought requires. In *Problems of Rationality* (pp. 135-150). Oxford: Oxford University Press.
- Davidson, D. (2004d). Paradoxes of irrationality. In *Problems of Rationality* (pp. 169-188). Oxford: Oxford University Press.
- Davidson, D. (2004e). Incoherence and irrationality. In *Problems of Rationality* (pp. 189-198). Oxford: Oxford University Press.

- Davidson, D. (2004f). Deception and division. In *Problems of Rationality* (pp. 199-212). Oxford: Oxford University Press.
- Davidson, D. (2004g). Representation and interpretation. In *Problems of Rationality* (pp. 87-100). Oxford: Oxford University Press.
- Davidson, D. (2004h). Who is fooled? In *Problems of Rationality* (pp. 213-230). Oxford: Oxford University Press.
- Dennett, D. (1987) *The Intentional Stance*. Cambridge, MA.: MIT Press.
- Dennett, D. (1998). Producing future by telling stories. In *Brainchildren: Essays on Designing Minds*. Cambridge, MA: MIT Press..
- Devitt, M. (1993). A critique of the case for semantic holism, *Philosophical Perspectives*, Vol. 7, Language and Logic, pp. 281-306.
- Devitt, M. (1995). *Coming to Our Senses: A Naturalistic Program for Semantic Localism*. Cambridge, UK: Cambridge University Press.
- Dragani, A. F., & Nissan, E. (2004). Salvaging the spirit of the meter-models tradition: a model of belief revision by way of an abstract idealization of response to incoming evidence delivery during the construction of proof in court. *Applied Artificial Intelligence* 18, 277-303.
- Duhem, P. (1954). *The Aim and Structure of Physical Theory*, translated by Phillip Wiener, Princeton, NJ: Princeton University Press.
- Dummett, M. (1993). Truth and meaning, in *The Seas of Language*. Oxford: Oxford University Press.
- Evnine, S. J. (1999). Believing conjunctions. *Synthese* 118, 201-227.
- Field, H. (1977). Logic, meaning and conceptual role, *Journal of Philosophy* 74, 379-409.
- Flanagan, O. (1988, April). Pragmatism, ethics, and correspondence truth: response to Gibson and Quine. *Ethics*, 98:3, 541-549.
- Fodor, J. (1983) *The Modularity of Mind*, Cambridge, MA: MIT Press.
- Fodor, J. (2000). *The Mind Doesn't Work that Way*. Cambridge, MA: MIT Press.
- Fodor, J. and Lepore, E. (1992). *Holism: A Shopper's Guide*. Oxford: Blackwell.
- Fodor, J. and Lepore, E. (eds.). (1993). *Holism: A Consumers Update*. Amsterdam: Rodopi.

- Foley, Richard. (1979). Justified inconsistent beliefs. *American Philosophical Quarterly* 16, pp. 247-57.
- Gibbons, J. (2001). Externalism and knowledge of the attitudes. *The Philosophical Quarterly*, 51:203, 13-28.
- Gillies, A. (2004). New foundations for epistemic change. *Synthese* 138, 1-48.
- Gillies, A. (2006). What might be the case after a change in view. *Journal of Philosophical Logic* 35, 117-145.
- Glüer, K. (2001). Wittgenstein and Davidson on agreement in judgment", in G. Oliveri (ed.) *From the Tractatus to the Tractatus and other Essays* (pp. 81-103). Frankfurt: Peter Lang.
- Guinn, C. I. (1998). An analysis of initiative selection in collaborative task-oriented discourse. *User Modeling and User-Adapted Interaction* 8:3-4, 255-314.
- Guttenplan, S. (1994). Holism, in S. Guttenplan (ed). *A Companion to the Philosophy of Mind*, Cambridge, UK: Blackwell.
- Haack, S. (2003). *Defending Science Within Reason*. Amherst, NY: Prometheus Books.
- Hansson, S. O. (2003). Ten philosophical problems in belief revision. *Journal of Logic and Computation*, 13:1, 37-49.
- Harman, G. (1968). Enumerative induction as inference to the best explanation. *The Journal of Philosophy*, 65:18, 529-533.
- Harman, G. (1974). *Thought*, Princeton, NJ: Princeton University Press.
- Harris, P. L. (2001). The veridicality assumption. *Mind & Language*, 16:3, 247-262.
- Hawley, P. (2003). Skepticism and the value of knowledge. (unpublished)
- Heil, J. (1989). Minds divided. *Mind*, 98:392, pp. 571-583.
- Hempel, C. (1950). Problems and changes in the empiricist criterion of meaning. *Revue internationale de Philosophie*, 4, 41-63.
- Hempel, C. (2000). Memoirs. In *Selected Philosophical Essays* (pp. 251-304). Cambridge, UK: Cambridge University Press.
- Hill, C.S. and Schecter, J. (2007). Hawthorne's lottery puzzle and the nature of belief, *Philosophical Issues*, 17:1, pp. 102-122.
- Hintikka, J. (1962). *Knowledge and Belief*. Ithaca, NY: Cornell University Press.

- Humphreys, P. (1989). *The Chances of Explanation*. Princeton, NJ: Princeton University Press.
- Jackman, H. (1999). Moderate holism and the instability thesis, *American Philosophical Quarterly*, 36:4, pp. 361-69.
- James, W. (1890.1950). *The Principles of Psychology*, v. 1, New York: Dover.
- Jones, T., E. Mulaire, and S. Stich. (1991). Staving off catastrophe: a critical notice of Jerry Fodor's Psychosemantics, *Mind and Language*, 6:58-82.
- Joyce, J. (unpublished). The accuracy of partial beliefs, I and II. Read at the Formal Epistemology Workshop at the University of California Berkeley, May 21-23, 2004.
- Jones, T. (2003). The virtues of nonreduction, even when reduction is a virtue. *The Philosophical Forum*, 34:2, 121-140.
- Kalderon, M.E. (2001). Reasoning and representing, *Philosophical Studies* 105:2, 2001, pp. 129-160.
- Kornblith, H. (1993). *Inductive Inference and Its Natural Ground*. Cambridge, MA: MIT University Press.
- Lepore, E. and Ludwig, K. (2005). *Donald Davidson: Meaning, Truth, Language, and Reality*, Oxford: Oxford University Press.
- Lewis, D. (1970). How to define theoretical terms. *The Journal of Philosophy* 67, pp. 427-446.
- Lewis, D. (1972). Psychophysical and theoretical identifications, *Australasian Journal of Philosophy*, 50:3, pp. 249 – 258.
- Lewis, D. (1982). Logic for equivocators, *Noûs*, 16:3, pp. 431-441.
- Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy*, 61:4, pp. 343-377.
- Lewis, D. (1986). *On the Plurality of Worlds*. Malden, MA: Blackwell.
- Lewis, D. (1999). Reduction of Mind, in *Papers in Metaphysics and Epistemology, Volume Two* (pp. 291-324). Cambridge, UK: Cambridge University Press.
- Lopes, D. (1997). Art media and the sense modalities: tactile pictures. *Philosophical Quarterly* 47, 425-440.
- Lormand, E. (1996). How to be a meaning holist, *The Journal of Philosophy*, Vol. XCIII:2, 51-73.

- Luper-Foy, S. (1990). Arbitrary reasons. In M.D. Roth and G. Ross (eds.) *Doubting*, Dordrecht: Kluwer Academic Publishers, pp. 39-55.
- MacKay, D. M. (1969). *Information, Mechanism and Meaning*. Cambridge, MA: MIT Press.
- Mandik, P. & Clark, A. (2002). Selective representing and world making. *Minds and Machines* 12:3, pp. 383-395.
- McCarthy, J., Hayes, P. J. (1969). Some philosophical problems from the standpoint of artificial intelligence. *Machine Intelligence*, 4, pp. 463-469.
- Mele, A. (1987). *Irrationality: An Essay on Akrasia, Self-Deception, and Self-Control*. Oxford: Oxford University Press.
- Misak, C. (1987). Pierce, Levi, and the aims of inquiry. *Philosophy of Science*, 54:2, pp. 256-265.
- Mizzaro, S. (2001). Towards a theory of epistemic information. in *Information Modelling and Knowledge Bases XII*. Amsterdam: IOS Press. pp. 1-20.
- Morganti, M. (2004). On the preferability of epistemic structural realism. *Synthese*, 142, pp. 81-107.
- Nagel, T. (1974). What is it like to be a bat?, *The Philosophical Review*, 83:4, pp. 435-50.
- Neale, S. (2001). Meaning, truth and ontology, in Kotatko, Pagin and Segal (eds.) *Interpreting Davidson* (pp. 155-197), Stanford: CSLI Publications.
- Nelson, K. & Nelson, R. R. (2002). On the nature and evolution of human know-how. *Research Policy*, 32, 719-733.
- Pagin, P. (2006). Meaning holism, in Ernie Lepore and Barry Smith (eds), *Handbook of Philosophy of Language* (pp. 213-232), Oxford: Oxford University Press.
- Papineau, D. (2003). Reply to Kirk and Melnyk. *Sito Web Italiano per la Filosofia Philosophy of Mind*, May 9 <http://www.swif.uniba.it/lei/mind/forums/papineau.htm>.
- Pears, D. (1984), *Motivated Irrationality*, Oxford: Clarendon Press.
- Pinker, S. (1994). *The Language Instinct: How the Mind Creates Language*. New York: Harper Collins.
- Pollock, J. (1997a). Reasoning about change and persistence: a solution to the frame problem. *Noûs*, 31:2, pp. 143-169.
- Pollock, J. (1997b). Taking perception seriously. *Proceedings of the First International Conference on Autonomous Agents*, Marina del Rey, California, Feb. 5-8, 1997, New York: ACM Press. pp. 526 – 527.

- Pollock, J. (1999). Rational Cognition in OSCAR. Lecture Notes In Computer Science; Vol. 1757 archive 6th International Workshop on Intelligent Agents VI, Agent Theories, Architectures, and Languages (ATAL), pp. 71-90.
- Putnam, H. (1998). *Representation and Reality*. Cambridge, MA: MIT Press.
- Quine, W. V. O. (1951). Two dogmas of empiricism. *Philosophical Review* 60, pp. 20-43.
- Quine, W. V. O. (1969). *Ontological Relativity and Other Essays*. New York: Columbia University Press.
- Quine, W.V.O. (1970). *The Web of Belief*. New York: Random House.
- Ram, A. (1993). Creative conceptual change. *Fifteenth Annual Conference of the Cognitive Science Society*, pp. 17-26.
- Ramsey, F. P. (1990). General propositions and causality, in D.H. Mellor (ed.) *Philosophical Papers* (pp. 145-63). Cambridge, UK: Cambridge University Press. pp. 145-63.
- Ramsey, F. P. (1931). Theories, in Frank Plumpton Ramsey, *The Foundations of Mathematics and Other Logical Essays*, edited R. B. Braithwaite. London: Kegan Paul.
- Ramsey, W., Stich, S. & Garon, J. (1990). Connectionism, eliminativism and the future of folk psychology, *Philosophical Perspectives*, Vol. 4, Action Theory and Philosophy of Mind, pp. 499-533.
- Rescher, N. (1979) *Cognitive Systematization*. Oxford: Basil Blackwell.
- Rescher, N. (2003). *Epistemology: An Introduction to the Theory of Knowledge*. New York: SUNY Press.
- Riggs, W. (unpublished). Insight, open-mindedness, and the highest epistemic good.
- Robbins, P. (2004.) To structure or not to structure. *Synthese* 139, pp. 55-80.
- Roorda, J. (1997). Fallibilism, ambivalence, and belief. *The Journal of Philosophy* 94, pp. 126-155.
- Rosenberg, J. F. (1981). On understanding the difficulty in understanding difficulty. In H. Parret & J. Bonveresse (eds.), *Meaning and Understanding* (pp. 29-43). Berlin and New York: Walter de Gruyter.
- Russell, B. (1900). *A Critical Exposition of the Philosophy of Leibniz*. Cambridge, UK: Cambridge University Press.
- Russell, B. (1985). *The Philosophy of Logical Atomism*, Peru, IL: Open Court.

- Sainsbury, R. M. (1986). Degree of belief and degrees of truth. *Philosophical Papers* 15, pp. 97-106.
- Schaffer, J. (2007). From nihilism to monism. *Australasian Journal of Philosophy* 85, pp. 175-91.
- Schaffer, J. (forthcoming a). The least discerning and most promiscuous truthmaker, forthcoming in *Philosophical Quarterly*.
- Schaffer, J. (forthcoming b). Monism: The priority of the whole, forthcoming in *Philosophical Review*.
- Schaffer, J. (in progress). The internal relatedness of all things.
- Seager, W. (2001). The constructed and the secret self. In A. Brook & R. DeVidi (eds.), *Self-Reference and Self-Awareness* (pp. 247-268). Amsterdam: John Benjamins.
- Seager, W. (unpublished). Emergence and efficacy.
- Sellars, W. (1963). Some reflections on language games, in Sellars *Science, Perception and Reality*, London: Routledge and Kegan Paul.
- Shoemaker, S. (1996). *The First-Person Perspective and Other Essays*. Cambridge, UK: Cambridge University Press.
- Smith, N., Hermelin, B., Tsimpli, I. (2003). Dissociation of social affect and theory of mind in a case of Asperger syndrome. *UCL Working Papers in Linguistics*, 15, 357-377.
- Soames, S. (1987). Direct reference, propositional attitudes and semantic content. *Philosophical Topics* 15: 47-87.
- Sosa, E. (2003). Knowledge of self, others and world. In K. Ludwig (ed.), *Contemporary Philosophy in Focus: Donald Davidson* (pp. 163-182). Cambridge, UK: Cambridge University Press.
- Sosa, E. (forthcoming). A defense of the use of intuitions in philosophy. *Stich and His Critics*, ed. Michael Bishop and Dominic Murphy, Oxford: Blackwell.
- Sperber, D. (1994). The modularity of thought and the epidemiology of representations. In L. Hirschfield & S. Gelman (eds.), *Mapping the Mind: Domain Specificity in Cognition and Culture* (pp. 39-67). Cambridge, UK: Cambridge University Press.
- Sperber, D. & Wilson, D. (1995). *Relevance: Communication and Cognition* (2nd ed.). Oxford: Blackwell.
- Stalnaker, R. (1984). *Inquiry*. Cambridge, MA: MIT Press.

- Stalnaker, R. (2001). Robert Stalnaker. In S. Guttenplan (ed.) *A Companion to the Philosophy of Mind* (pp. 561-568). Oxford: Blackwell.
- Stalnaker, R. (2004). Lewis on intentionality. *Australasian Journal of Philosophy*, 82:1, pp. 199-212.
- Stalnaker, R. (2006). Assertion. In A. P. Martinich (ed.) *The Philosophy of Language*, 5th edition, Oxford: Oxford University Press.
- Stich, S. (1983). *From Folk Psychology to Cognitive Science: The Case Against Belief*. Cambridge, MA: MIT Press.
- Tooby, J., & Cosmides, L. (1995). Foreword. In S. Baron-Cohen, *Mindblindness: An Essay on Autism and Theory of Mind* (xi-xviii). Cambridge, MA: MIT Press.
- Tye, M. (2003). *Consciousness and Persons: Unity and Identity*. Cambridge, MA: MIT Press.
- van Fraassen, B. (1985). Replies. In P. M. Churchland & C. A. Hooker, *Images of Science: Essays on Realism and Empiricism* (pp. 245-308). Chicago: University of Chicago Press.
- van Fraassen, B. (1994). Science, probability and the proposition. *Proceedings of the Biennial Meeting of the Philosophy of Science Association* 2, pp. 339-348.
- van Fraassen, B. (1995, August). Fine-grained opinion, probability, and the logic of full belief. *Journal of Philosophical Logic* 24, pp. 349-377.
- van Fraassen, B. (2001). Constructive empiricism now. *Philosophical Studies* 106, pp. 151-170.
- Van Gulick, R. (1993). "Who's in Charge Here? And Who's Doing All the Work?". In J. Heil & A. Mele (eds.), *Mental Causation* (pp. 233-256). Oxford: Oxford University Press.
- van Inwagen, P. (1987) When are objects parts? *Philosophical Perspectives*, Vol. 1, Metaphysics, pp. 21-47.
- van Inwagen, P. (1990) *Material Beings*. Ithaca, NY: Cornell University Press.
- Varzi, A. (2003). Mereology, in *Stanford Encyclopedia of Philosophy*.
- Vermazen, B. The intelligibility of massive error. *The Philosophical Quarterly*, 33:130, pp. 69-74.
- Vogeley, K., Kurthen, M., Falkai, P., & Maier, W. (1999). Essential functions of the human self model are implemented in the prefrontal cortex. *Consciousness and Cognition*, 8, pp. 343-393.
- Waskan, J. & Bechtel, W. (1997). Directions in Connectionist Research: Tractable Computations without Syntactically Structured Representations. *Metaphilosophy*, 28:1 & 2, pp. 31-62.

- Wedin, M. (2004). Aristotle on the firmness of the principle of non-contradiction. *Phronesis* XLIX 2004, pp. 225-265.
- Williams, M. (2001). *Problems of Knowledge*. Oxford: Oxford University Press.
- Wilson, N. L. (1959). Substances without substrata, *Review of Metaphysics* 12, pp. 521-39.
- Wilson, N. L. (1967). Linguistic butter and philosophical parsnips. *The Journal of Philosophy*, 64:2, 55-67.
- Wittgenstein, L. (1953). *Philosophical Investigations* (G. E. M. Anscombe, Trans.), Oxford: Basil Blackwell.
- Wittgenstein, L. (1961). *Tractatus logico-philosophicus*, translated D.F. Pears and B.F. McGuinness, London: Routledge and Kegan Paul.
- Wollheim, R. (1999.) *On The Emotions*. New Haven, CT: Yale University Press.
- Zeevat, H. (2003). Particles: presupposition triggers, context markers or speech act markers, in *Optimality Theory and Pragmatics* (pp. 91-111). R. Blutner & H. Zeevat (eds), Basingstoke: Palgrave Macmillan.