Patterns of Perception

Nemira Hathor Gasiunas

Submitted in partial fulfilment of the requirements

for the degree of Doctor of Philosophy

in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2019

ABSTRACT

Patterns of Perception

Nemira Gasiunas

Our ordinary concept of perception contains a seeming tension: we distinguish perception from thought on the grounds that it is a *direct* awareness of mind-independent objects through their effects on our senses; yet we also allow that *what we see (hear, feel, etc)* is determined by how we interpret or classify the data that comes through our senses. Theorists of perception disagree over which of these intuitions should prevail, with some maintaining that concepts are in play all the way down and others that perceptual awareness is wholly immediate and concrete. But we do not have to choose. This dissertation argues that the patterns of perception sustain a distinctive form of nonconceptual classification, in which *property spaces* organize sensory matter so as to preserve rather than discard its concreteness and detail. What then is classification without concepts? What sort of abstraction, generality, representation, or form does it entail? And what ramifications then for thinking about the roots of language and reason, and of our awareness of the external world?

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGEMENTS

The writing of this dissertation challenged me in ways I could not have imagined. So, whilst authors often acknowledge those 'without whom this work would not have been written,' I want to stress how literally these words are true for me.

At a time when I had nothing but a head full of ideas and a bunch of hang-ups to show for five years of near-constant toil, Lydia Goehr took me on as her student. If I am sitting here not 14 months later putting the finishing touches on a 300-page text, she is the reason why. Lydia believed in me and she fought for me when I had almost no fight left in me. She reignited my love for the written word, and taught me by her example to see that my creativity and intuition are an asset to my philosophical abilities and not a weakness in them. If I can give to my future students even a fraction of what was so freely given to me, I will have done well by them.

If writing can be a hell, then my friendship with Manu Chander is steel forged in its fires: his cheers propelled me over the finish line, just as loud as they had been when the starting gun sounded seven years ago. The same goes for Jorge Morales, who walked this path beside me and understands those things I do not know how to communicate in words. Matthew Heeney was wise and impossibly generous in coaxing the real me out from underneath an armour of words. I do not know if he, Simon Brown, Grace Helton and Louise Doust realize how important their lively discussions and compassionate feedback were in making philosophy feel fun — and safe — for me again. Lex Braes kept my eyes and mind open through his intellectual honesty and his passion for perceptual experience. David Rosenthal, through his encouragement and endless wisdom, has helped me in ways I could not have imagined when I first took his class as part of my masters requirements

iv

Glory be to God for dappled things—
For skies of couple-colour as a brindled cow;
For rose-moles all in stipple upon trout that swim;
Fresh-firecoal chestnut-falls; finches' wings;
Landscape plotted and pieced—fold, fallow, and plough;
And all trades, their gear and tackle and trim.

All things counter, original, spare, strange;
Whatever is fickle, freckled (who knows how?)
With swift, slow; sweet, sour; adazzle, dim;
He fathers-forth whose beauty is past change:
                    Praise Him.

                    —Gerard Manley-Hopkins, 'Pied Beauty'

INTRODUCTION

Our ordinary concept of perception contains two elements which, if not quite contradictory, stand in a seeming tension with one another. In the first place, we regard perception as yielding a *direct* awareness of objects in our environment, through their effects on our senses. If you see a bird in the tree, then you are aware of an entity, existing independently of you that is directly registered by your visual apparatus. By contrast, if you conclude that there is a bird in the tree on the basis of seeing its droppings on the ground, or if you are told there is a bird in the tree by someone else who can see it, then you apprehend the existence of the bird *indirectly*; not by perceiving it but by constructing it out of something else you perceive. Yet we also readily allow that what we see (hear, feel, etc) is determined by how we interpret or classify the data that comes through our senses. In his well-loved children's story, *The Little Prince*, Antoine de St Exupery illustrates this idea by recounting his early disappointment in the world of adults. Having been inspired to draw a boa-constrictor digesting an elephant, his younger self proudly shows his drawing to the adults—only to find they all mistake his drawing for a hat.
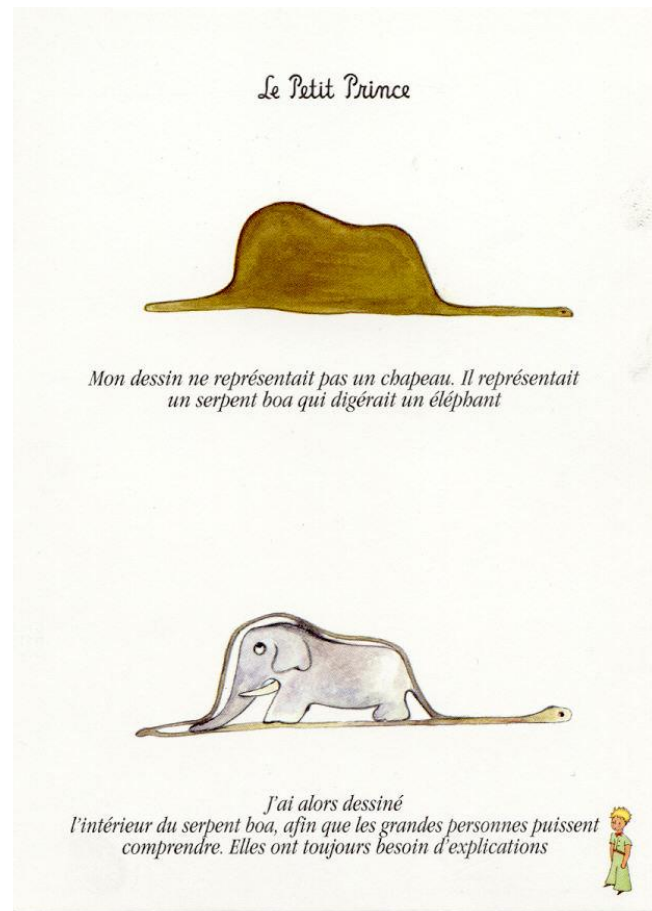
Figure 1. Two drawings of a boa-constrictor having swallowed an elephant. The first figure illustrates the author's drawing in its original form, which is subsequently mistaken by the adults for a drawing of a hat. In the second figure, the snake is drawn in cross-section, so that the true meaning of his figure is impossible to misconstrue, even by the stupidity of adults.

If it is possible for two people to be presented with the same data and yet find an entirely different objective significance in it, then seeing something as a bird or as a hat or as an elephant-engorged boa constrictor cannot simply be a matter of taking in what is there. Rather, how we perceive the world to be must depend as much on facts about how we ourselves are constituted—on how we interpret or classify the data that comes through our senses, and on the body of assumptions and expectations we bring to bear in making those

interpretations—as on facts about the world.[1] The question is how to reconcile this idea with the intuition that perception is a distinctively direct mode of access to objects and their properties. Notice that the question still arises even if we focus our attention on basic perceptual magnitudes: the perception of an object with a brown color and a 'hat-like' form, on which such classifications as 'hat' or 'elephant-engorged boa-constrictor' would seem to be based, are formed on the basis of retinal images which are *multivalent*, that is, capable of being read in more than one way. The fact that we sometimes get the world wrong does not mean, of course, that we do not mostly get it right. Yet the phenomenon of conflicting interpretations does seem to threaten our intuitive ideas about what is special to perceptual objectivity.

Theoretical discussions of perception, which implicitly accept this tension, are divided as to which of these two intuitions should prevail. One tradition, which includes Gottlob Frege (1918/1956), John McDowell (1994, 1994a, 2009), Alex Byrne (2005) and Jerry Fodor (1975, 2015) embraces the indirect, interpretative aspect of experience. In the view of these *conceptualists*, perceptual experience of even the most basic properties is a matter of subsuming objects under concepts. Another tradition, which includes Charles Travis (2004, 2007), Fred Dretske (1981), and Tim Crane (1992, 2009, 2012), embraces the directness of perception. In the view of these *naive realists*, perception is not a classification, perhaps not even a form of interpretation or representation, but is rather an

---

[1] Indeed, these observations explain why the narrator, as an adult, carries his childhood drawing around with him as a test of character: "Whenever I encountered a grown-up who seemed to me at all enlightened, I would experiment on him with my drawing Number One, which I have always kept. I wanted to see if he really understood anything. But he would always answer, "That's a hat." Then I wouldn't talk about boa constrictors or jungles or stars. I would put myself on his level and talk about bridge and golf and politics and neckties. And my grown-up was glad to know such a reasonable person" (2000: 3).

awareness of an immediate and concrete kind. But whilst these two views each get something importantly right, they also miss something, and it is my aim to fill in that gap. What if there were a middle way between these views? What if there were a way to capture both the idea that perception must classify the world, and that we cannot regard it as doing so in the manner of conception, that is through 'leaving out' the particularity of sense?

In this dissertation, I shall argue that the patterns of perception, which organize, interpret and process the matter of the senses in a nonconceptual way, are *both* genuinely direct *and* inherently abstract. The patterns of perception stand alongside the patterns of thought as two fundamentally distinct organizing principles or systems. In taking something as a hat or as a snake-engorged elephant, we interpret cognitively and conceptually. Patterns of thought subsume a perception of the object, picked out through its color, form, location and so forth, under an atomic or word-like representation of a category. This form of classification is facilitated by a thinker's capacities to draw inferences, through knowledge of law-like relations, about the likely causes of their perceptions. In perceiving an object's color, form, position, or location, meanwhile we are also classifying and interpreting — but not through concepts. Patterns of perception classify objects as being of some type by placing the sensations caused by those objects within 'property spaces' or holistic frames of reference, which organize sensory matter so as to preserve rather than discard its concreteness and detail. This form of classification is facilitated by patterns of predictive processing, in which the structures of perceptual domains are figured out in advance. Armed with this understanding of perceptual

representation, we do not have to choose between one or other of the two facets of perception. The tension dissolves.

To better understand what has led many philosophers to mistakenly regard classification as inherently opposed to concreteness, it is helpful to look briefly at a classical view of beauty, as offered by Plotinus:

> All shapelessness whose kind admits of pattern and form, as long as it remains outside of Reason and Idea, is ugly from that very isolation from the Divine-Thought. And this is the Absolute Ugly: an ugly thing is something that has not been entirely mastered by pattern, that is by Reason, the Matter not yielding at all points and in all respects to Ideal-Form. But where the Ideal-Form has entered, it has grouped and coordinated what from a diversity of parts was to become a unity: it has rallied confusion into co-operation: it has made the sum one harmonious coherence: for the Idea is a unity and what it moulds must come into unity as far as multiplicity may. (trans. 2018: I.6.1)

What Plotinus discusses under the name of unity is what contemporary philosophers call classification: it is a mental act or capacity by which qualitatively distinct individuals are 'grouped and coordinated' under a single common type. To speak of unity as something exclusive of variation within a kind, meanwhile, is to describe the way in which *concepts* classify. Concepts work so as to exclude difference through a pattern that homogenizes. They achieve unity by "discarding ... individual differences and by forgetting the

distinguishing aspects … by overlooking what is individual and actual" (Nietzsche, 1873/2006: 117). It is this assumption that conceptual classification is definitive of unity— that subjects can group and coordinate *only* through discarding—that leads to the polarization I am seeking to challenge. Those in the conceptualist tradition think perception must be conceptual in part *because they think perception must classify objects*. Those in the 'direct' tradition deny that perception classifies the world in any form *because they deny* that perception classifies the world akin to conceptual thought. Plotinus's aesthetic orientation is also suggestive for understanding the problem with contemporary approaches in a further way. We might put the point like this: since the two parties are engaged in debate with one another, they are obviously aware of the facts which motivate the other side. Each finds their own motivations to be more compelling. And yet the motivations for the other side have not gone away. In *The Death and Life of Great American Cities*, Jane Jacobs criticizes urban planners for having been blinded by a form of intellectual arrogance which ignores facts about cities which are readily observable to anyone prepared to spend a moment walking around in them. What takes them on this path, she says, is not a conscious disdain for knowing how things work, but an over-reliance on what "the saints and sages of modern orthodox planning have said about how cities *ought* to work and what *ought* to be good for the people and businesses in them." For their devotion to these theoretical dictates is such that "when contradictory reality intrudes, threatening to shatter their dearly won learning, they must shrug reality aside" (1961: 8). It is likewise, I wish to suggest, with theorists of perception, inasmuch as they have continued to insist on an assumption which precludes our capturing the *dual* aspect of perception. If each has

clung to their theoretical presuppositions, it is for having failed to look; or better, for having failed to see. For the dual nature of perception is not hidden: it is already open to view.

We find (to continue this aesthetic theme) a rather different vision of what is worth valuing in Gerard Manley-Hopkins, whom I quoted above: namely, an "aesthetic realism". As a critic describes this approach,

> With Hopkins, the religious way of seeing welcomed precision, detail, perceptive speed, the determination of courageous observing. Aesthetic Realism stands for the utter regard of any object whatsoever—in all its ontological nonconformity—and the deepest respect for existence, as such, as the cause of value. (Siegel, 1981: 34)

Hopkins' eulogizing of dappled things (the variety of god's creations, mind-independent objects) is also a eulogizing of our dappled perceptions of things. To delight in 'fresh-firecoal chestnut-falls' is also to delight in the richness of the visual experience of these colors and textures: the fallen chestnuts resemble coals bursting in a fire because of how the chestnuts' reddish-brown meat is exposed when the shells break against the ground. The image is of the landscape which has been "plotted and pieced" into fields (like quilt squares) by agriculture. In short, one cannot sing in praise of nature, in all her rich variety, without singing in praise of the perceptual capacities as our mode access to this richness. In identifying the atomistic abstraction embodied in concepts, words, and forms, Nietzsche's point was to highlight the sense in which these forms of representation are 'lies.' For

although they purport to *represent* nature, "nature is acquainted with no forms and no concepts, and likewise with no species, but only with an X which remains inaccessible and undefinable for us." By their discarding of particularity, they never touch the reality of the thing, the thing *as it is in itself:* "we believe that we know something about the things themselves when we speak of trees, colors, snow, and flowers; and yet we possess nothing but metaphors for things—metaphors which correspond in no way to the original entities." My point, in this dissertation, is that the reason perceptual experience is capable of serving as our primary link to the world, the foundation of our thought about things, is precisely because of the reality it preserves within its representations. But—and here is the other, equally important point—the respect for the object in all its detail means recognizing all of it, including what may seem from the point of view of our theoretical assumptions 'counter, spare, original, strange': that this particularity figures in conjunction with classification. Perception both abstracts from particularity and preserves it: a surface seen at different distances or under different illuminants looks *both* the same (in its objective features) and different (in the specificity of how it strikes us). This dissertation, inasmuch as it calls for an acknowledgement of the varieties of representation, sings in praise of the dappled forms of representation: the varying ways in which creatures organize, classify, make sense of their environments.

Not the least of the advantages to accrue from studying perception on its own terms, and not through the lens of thought, is a more accurate view of the power of thought itself. Arguments for the existence of a 'nonconceptual content' in perception are often presented as a reaction against the impulse to over-intellectualize perception, but I think it would be

equally accurate to say (at least of the present account) that it is a reaction against an impulse to 'under-intellectualize' thought. When we elevate perception to the status of classification without yet subsuming this under conceptual classification and conceptual objectivity, we allow ourselves to see what is truly remarkable in conceptual classification: not abstraction per se, but a special form of abstraction which consists in the ability to disconnect concepts from a context, to atomize them. In this lies the source of truly conceptual generality; the roots of language and reason. Thus, by acknowledging the existence of nonconceptual content we open ourselves to an appreciation of what is special in conception as much as in perception.

***

We cannot consider the question, 'does perception have a nonconceptual content?' without engaging with three significant areas of debate within the philosophy of mind. One is the subject of 'the nonconceptual' as such. How shall we give substance to the idea of representation without concepts? If, indeed, we can make sense of this at all? A second topic concerns the nature of mental representation. What does it mean to say that a state 'has' a representational content? What are the criteria by which such attributions are justified? The third topic concerns the nature of perception itself. To ask whether perception has a different kind of content from thought is to assume perception cannot be reduced to thought. But what is it to perceive the world—as opposed to having beliefs about it? Although in principle, it is possible to consider each of these questions independently, one of the things I hope to convey in this dissertation is the degree to which

these questions form part of an interconnected picture. Indeed, it is this holistic character of the theory of content that is in large part responsible for making the 'nonconceptualism debate' as contentious and difficult as it is.

1) The Nonconceptual. As Hubert Dreyfus has observed, there have always been those— "painters, writers, historians, linguists, philosophers in the romantic tradition, Wittgensteinians, and existential phenomenologists—who have felt there is another kind of intelligibility which gets us in touch with reality besides the conceptual kind" (2001: xv). In this sense, the 'nonconceptualism debate' concerning perception is part of a larger disagreement over the possibility of representation outside of the paradigm of conceptual thought. Indeed, we may get a feel for the recent debate over perception by considering a disagreement between P.F. Strawson and Donald Davidson which in many ways lies at its origins. Consider the opening remarks of Strawson's *Individuals*:

> We think of the world as containing particular things, some of which are independent of ourselves; we think of the world's history as made up of particular episodes in which we may or may not have a part and we think of these particular things and events as included in the topics by common discourse as things about which we can talk to each other. These are remarks about the way we think of the world, about our conceptual scheme. (1959: 15)

Begin by noticing a surprising aspect to Strawson's claims. His ostensible subject matter is the subject-matter of metaphysics. The particulars and episodes we think of the world as containing are things we think the world contains *independent of our thought about them*; part of the basic structure of a sharable, mind-independent reality. It is therefore surprising when Strawson frames these observations as 'remarks about our conceptual scheme.' At one level, it is possible to see *Individuals* as an exploration of reference; the phenomenon by which our thoughts, and speech, come to be directed onto entities in the world. But there is also another connotation: that any analysis of metaphysical structure will also, and inevitably, be an analysis of how we think—because our perspective on mind-independent reality is determined by *our system of representation*, that is, by the concepts and 'mechanisms of objectivity' we bring to bear in thinking and speaking about that reality. Davidson's opening remarks, in 'On the Very Idea of a Conceptual Scheme,' speak directly to this idea:

> There may be no translating from one scheme to another, in which case the beliefs, desires, hopes and bits of knowledge which characterize one person have no true counterparts for the subscriber to another scheme. Reality itself is relative to a scheme: what counts as real in one system may not in another.
>
> Even those thinkers who are certain there is only one conceptual scheme are in the sway of the scheme concept; even monotheists have religion. And when someone sets out to describe "our conceptual scheme," his homey task assumes, if we take him literally, that there might be rival systems.

Conceptual relativism is a heady and exotic doctrine, or would be if we could make good sense of it. The trouble is, as so often in philosophy, it is hard to improve intelligibility while retaining the excitement. (1973: 5)

When Strawson speaks of 'our conceptual scheme,' he commits himself to what Davidson calls 'the dualism of scheme and content,' that is, he assumes a sharp separation between an organizing structure and a data to be organized. Strawson is also, therefore, committed to an idea which goes hand in hand with this dualism: the possibility that our conceptual scheme is only one of many equally valid ways of organizing reality, incommensurable with our own. He is committed to conceptual relativism. This aspect of Strawson's view is in evidence, for example, in the second chapter of *Individuals* which explores a world of imaginary creatures who only have capacities for hearing, and tries to determine what kinds of thoughts the inhabitants of this 'sound world,' would be capable of having. One can certainly see why Davidson calls relativism a "heady and exotic doctrine." Yet he regards the dualism of scheme and content to which it commits us as a fundamentally incoherent position: "Different points of view make sense, but only if there is a common coordinate system on which to plot them; yet the existence of a common system belies the claim of dramatic incomparability" (ibid., 6). We can readily acknowledge 'modest' differences between systems of representation: a speaker of French organizes the same subject matter differently than the speaker of Russian; a map presents a landscape differently from a verbal description of that same landscape. But it is clear that the conceptual relativist intends something deeper than these contrasts. They are talking about a difference at the

level of thought-content; that which is evaluable as true or false with respect to how the world is. Davidson thinks the relativist cannot give an account of how a representational content performs that function in an 'incommensurably different' way, without collapsing into incoherence or paradox. The form of our thought about the world is not something in principle separable from its being a thought *about the world* at all.

There are important parallels between the project Strawson was engaged in and that of Gareth Evans, who was his student. Not only does Evans's *Varieties of Reference* take up the theme of reference-based thought and language; it also argues that an adequate explanation of this phenomenon requires us to recognize a 'nonconceptual content' for perception (indeed it is Evans who is widely credited with introducing the term into the contemporary philosophical discourse). Although a 'nonconceptual content' is, by definition, not a 'conceptual scheme,' the two ideas are naturally connected. For what is 'the nonconceptual content of perception' if not the 'ineffable': a way of organizing the world for which 'there may be no translating' into thought? And so, it is perhaps not surprising that we can draw a direct line between Davidson's rejection of the 'conceptual scheme' metaphor and those who have rejected the very idea of a nonconceptual content. In John McDowell's *Mind and World*, Davidson's dualism of 'scheme and content' is resurrected to serve as a warning to any who attempt to locate experience *outside* the conceptual realm— and so outside of the realm of what can meaningfully be said to constitute *reasons* for our thoughts about the world. Odd though it may seem, there is also a connection here with figures such as Charles Travis and Bill Brewer, who rejects the model of representation in favor of one on which perception is a direct relation to mind-independent objects. Despite

the fact that their view of experience is one Davidson would surely have rejected, it is he—along with other 'arch conceptualists' such as Frege—who paved the road that led them there. To invoke a *dualism* of scheme and content is to be presented with a binary choice: either experience is conceptual, or it is not a form of representation at all. And since it is their view that perceptual awareness *must be* outside of the conceptual realm, they conclude that perception must reside outside of the realm of representation. Similarly, McDowell opposes 'nonconceptual content' in large part *because* he regards it as a 'deceptive' or 'fraudulent' use of the label 'content.' For all their differences, then, Travis and McDowell share, and are partly motivated by, a common disagreement with the nonconceptualist that is of a different order of magnitude: there is no such thing as a 'nonconceptual content.' Representational content is essentially conceptual.

Earlier, I noted how Strawson explores the nature of thought through his 'descriptive metaphysics.' The theory of perceptual nonconceptual content I defend in this dissertation is likewise best understood in relation to a metaphysical observation, although one concerning properties rather than particulars. Some properties, amongst them length, weight, pitch and lightness, are *magnitudes*: continuous quantities which belong in ordered systems. A number of theorists, amongst them Brett Mundy (1988), Chris Swoyer (1987), and Christopher Peacocke (2014) have argued for a 'robust realism' about magnitudes, as irreducible elements in our ontology. This amounts to the claim that magnitudes are a special subclass of property, which stand in necessary *second order relations* to one another. In my view, we should regard perception as having a representational content which is organized in the structure of magnitudes. Like conceptual, propositional content,

this form of nonconceptual content is *structured*, containing both singular and attributive elements. But in nonconceptual content, the attributive element is not a predicate concept but a 'space' (or, better, a location in a space). In 'spatial attribution,' a sensory magnitude is represented as being of a certain type (for instance, a given spatial magnitude) by locating it in the part of the 'frame of reference,' which refers to that type.   This theory of nonconceptual content constitutes a true 'middle way' between the extremes of conceptual content and pure particularity. It challenges two false equivalencies which tend to animate opposition to nonconceptual content. One is to equate classification, the function by which objects are represented as being of the same *type*, with predicate concepts. The other is to equate representational content with propositions, fully abstract objects which represent states of affairs in context-independent ways. Perceptual classification is not predicative because the representation of any given piece of information is always accompanied by information concerning its relation to a *system* of properties. Moreover, given this form, a perceptual state functions to apply to particulars in a singular, context-dependent way. The form of attribution is non-propositional because it is object-involving; and because the objective content is always accompanied by a subjective aspect. It avoids absurdity by carving out a place for the nonconceptual *within* the space of classification. In contrast to those proponents of 'nonconceptual content' (including not only Evans but also Tim Crane (1992, 2009, 2012) and Fred Dretske (1982)) for whom this is something impoverished or 'pre-objective, the present account is a picture of a higher and lower forms of objectivity, marked by a difference in the degree to which representation of particulars depends on

features of the subject themselves. Despite Davidson's warning, then, nonconceptual content neither collapses into a 'merely modest' difference nor falls into incoherence.

2) Representational Content. We appeal to a notion of mental representation when we invoke a subject's thoughts about the world—what a subject believes, knows, remembers, intends, or desires—to explain their actions. The representational content of a thought specifies how the subject takes the world to be in terms of the conditions under which it would be true. For example, a natural explanation for why I bring my winter clothes out from storage as the weather starts to turn is that I believe winter is coming; where my having that belief is a matter of my bearing a relation to a representational content which is true if and only if winter is coming. These observations reflect the special status of thoughts as objective mental states, that is, mind-dependent entities which are essentially 'of' or 'about' mind-independent objects, properties, and events. By framing thought as a posture or stance which is correctly or incorrectly adopted according to how things are in the world, as McDowell puts it, we capture the idea that "thinking that aims at judgement, or at the fixation of belief, is answerable to the world—to how things are—for whether or not it is correctly executed" (1994: xii). Yet despite its centrality in our day to day lives, it is a difficult thing to say what exactly this idea of genuine psychological explanation consists in. Consider, for example, the fact that deciduous trees also 'prepare' for winter by dropping their leaves when the temperature gets consistently cold. We do not think this is because the trees *believe* winter is decisively coming, or that they predict consistently cold weather. The basis for distinguishing these cases seems clear enough: the attribution of

belief to a tree would be explanatorily idle; to mention a belief in connection with this process would be to project psychological properties onto a process that is already completely explained in *biological* terms. Yet the problem is to say *what it is* for a process to be 'genuinely explained' by attributions of thoughts or mental representations more generally. The problem becomes salient when one considers the many cases in which our intuitions are not clear-cut. Does the infant smile because they recognize their mother has returned or because they have gout? When I am packing for a trip and my dog, Walter, gets anxious, is this because he understands *that I am packing*, and 'knows' or 'believes' I am leaving for a trip—or is his reaction somehow more automatic, 'mindless'? Do the bluejays who skim the milk off the top of their milk-bottles reason (as my father believes) about the displacement of liquid? The further we drift further from the paradigm case of adult human activity, the less sure we are about how such questions could be decided.

Many philosophers have taken a complacent stance towards this problem, the tacit assumption being that our intuitions about explanatory relevance are reliable *enough*. The inadequacy of such an approach, particularly with regards to perception, may be illustrated by considering an example discussed by Michael Tye:

> A child as young as two months, upon tasting a little chocolate, typically behaves in a way that signifies that he/she wants more. The child will open and close its lips, push forward towards the chocolate, look happy. Why? The answer is that the chocolate tastes good. That's why the child wants more. The child's gustatory experience represents a certain taste and the child experiences that taste as good.

The taste is experienced as good by the child in that the child undergoes an overall experience which represents the presence of the taste in the mouth and represents it as good. Intuitively, this is not a cognitive response. It does not require its subject to possess evaluative concepts. (2005: 225)

Tye's claim that evaluative experiences have a nonconceptual content relies on the intuitive consideration that we need some notion of representational content to explain why the subject behaves one way rather than another (why the child acts in that way, and not in some other way); whilst we do not need to invoke the sophisticated capacities we associate with conceptual thought. But consider a rough way of working out an explanation for Tye's question in which the child's response can be fully accounted for without any appeal to *representation* at all. The child's taste-buds are causally sensitive to various chemical compound (or combination of compounds) in the chocolate: sugar, cocoa etc. The gustatory sensory system is causally sensitive to—carries information about—these types of proximal stimulation (Without receptors for a particular flavor, animals cannot be motivated to eat or avoid substances with that flavor). If sugar confers a selection benefit, humans could have evolved a functional association between the chemical compounds in sugar and an experience of 'craving' or pleasure, causing them to act in ways that will get them more sugar. For example, this might be due to the fact that for our primate ancestors, fruits, an energy-rich food source, were scarce in comparison to vegetables, and carries energy compared to fruits. Since the who consumed the most calories were the ones best able to stave off starvation and pass on their genes, and since fruits were relatively rare, the

primates who ate the most fruits were best positioned to win the battle of evolution. As it happens, I suspect the story I have just told—which is an entirely 'mindless' story, to which states with representational contents make no essential contribution—is probably closer to the truth than the one Tye gives us. But the present point is just that Tye owes us an argument that his proposal, rather than the one I just offered, is right. In the absence of reasons to suppose that the nature of the case is such that the child's behavior cannot be explained *except by* reverting to representational mental states, the argument offers us no reason for thinking that what we have is an instance of nonconceptual *content*, rather than something that is not a form of content at all.

I will argue for a robustly realist view of representational content, and likewise distinct *kinds* of representational contents, that is anchored in the notion of a psychological capacity or process that explains patterns of successful interactions with the environment. Living organisms have evolved a range of capacities allowing them to engage with their mind-independent environments in ways that fulfil their biological functions. Capacities for world-directed behavior need not, in themselves, involve representation. In the example above, I indicated how an action which is explained as a response to a delicious taste does not require us to assume a *representation of deliciousness* that is explaining the action. For the phenomenon can be captured in terms of the notion of an 'information registration' state (involving simple causal covariation), together with a biological function the correlation subserves. Following Tyler Burge, I regard representational states as tied to 'objectifying capacities': a *way of overcoming underdetermination* which involves systematically distinguishing 'the merely proximal from the probably environmental.'

(Burge, 2010). What is distinctive to objectifying capacities is that they are underwritten by processes in which states with representational contents figure both as causes and products. These capacities are tied to specific patterns of success and failure in engagement with environmental entities, properties and events.   Notice that the foregoing account is simultaneously an argument for thoughts having a representational content, and for their having a specific kind of content. This is an argument for representational content because these are entities which are invoked to explain how that works. As intermediaries. And it is also an argument for conceptual content in thought—for the same reasons. The role of representational content in marking distinctions between psychological kinds: these kinds are distinguished by *how* they acquire information about objects in the environment. I also offer a corresponding account of perception, and what distinguishes perceptual representation from thought (see following section).

Connecting the foregoing suggestions with my arguments concerning nonconceptual content, there is a story to be told here about the role of the a priori in theorizing about perceptual content, and mental content more generally. In arguing for the existence of a nonconceptual, non-propositional form of classification in perception, I strongly oppose a tendency to rule out the possibility of nonconceptual content on an a priori basis. Yet in the arena of psychological explanation, a role for a priori constraints on attributions of content is not merely allowable but essential. It is an a priori or necessary truth that in order to count as *representing* a given state of affairs, a subject must exhibit a certain pattern of general capacities to respond to that state of affairs; a pattern that would be meaningfully explained by that kind of state. Where there is no evidence of such

patterns, there is no evidence of a state that would causally explain these patterns. As we saw with Tye, a failure to impose these kinds of a priori constraints may result in erroneous conclusions about the nature of a creature's capacities to engage with their environment.

3) Perception. "We find certain aspects of seeing puzzling," Wittgenstein writes, "because we do not find the business of seeing puzzling enough." (1950/1977) The context for this comment is his description of an experience of looking a black-and-white photograph:

> I see in a photograph (not a colour photograph) a man with dark hair and a boy with slicked-back blond hair standing in front of a kind of lathe, which is made in part of castings painted black, and in part of smooth axles, gears, etc., and next to it a grating made of light galvanized wire. I see the finished iron surfaces as iron-coloured, the boy's hair as blond, the grating as zinc-coloured, despite the fact that everything is depicted in lighter and darker tones of the photographic paper. (ibid., 10)

He then proceeds to frame a question:

> But do I really see the hair blond in the photograph? And what can be said in favor of this? What reaction of the viewer is supposed to show that he sees the hair blond, and doesn't just conclude from the shades of the photograph that it is blond?—If I

were asked to describe the photograph I would do so in the most direct manner

with these words. If this way of describing it won't do, then I would have to start

looking for another. (ibid., 11)

Wittgenstein's point here is that, although we have firm intuitions that there is such a thing

as "really seeing" something, as opposed to making a judgement or drawing a conclusion

about what one sees, the question of what it is to 'genuinely see' a thing proves remarkably

difficult to answer. Suppose we say that the experience of 'blondeness' in the black and

white photograph does *not* count as a case of genuine seeing. It is true that if we take

Wittgenstein's description of his experience at face-value, he has a visual *experience* which

presents the color, blond, as immediately and non-inferentially present in the photograph

(specifically, in the part which represents the boy's hair) — but that experiential

'givenness' is not the relevant notion of givenness for this to count as a case of genuine

seeing. For there is another, more important sense in which the blondeness clearly is *not*

'present in' the photograph: the property, blonde, is a color; but the photograph is by

stipulation achromatic, *colorless*. For this reason, one could only have "gotten to" the

blondeness based on seeing the photograph via an inference; by drawing on prior

familiarity with the objects depicted and the nature of black-and-white photography.

Whatever genuine seeing is like, we feel, it is not like this. Yet a problem emerges when we

seek to make this idea precise. The idea of a perceiving subject as a 'recipient' of

information about the world presupposes the idea of what is perceived as being 'given' to a

subject. But whereas reflection on the phenomenology of perceptual experience

encourages a view on which the eyes take a 'snapshot' of what's there, when we probe this idea, we see that it cannot be taken literally. Perception is a representation: perceptual content is not 'in' the sensation; perceptual awareness comes from the mind. But once this abstraction is granted, it is hard to see how perception is really any different from applying a judgement.

The solution to this problem, I believe, lies in combining the account of nonconceptual content I have outlined with an understanding of how perception works as an autonomous system of representation. The difference between perception and thought is a matter of how they achieve their function of giving us knowledge about the world; the processes by which they produce objective representations of reality and which, most of the time, get reality right. What it is for something to be a thought, a belief or judgement, is for it to have been produced through a process of inference. What it is for something to be a perception is for it to be produced through a different kind of capacity. The view I propose may be understood against the background of two common approaches to perceptual processing. On one side are 'direct' theories of perceptual processing, most completely worked out in J.J. Gibson's 'ecological theory' of perception (1960; 1963; 2014), which preserve the naive view in a literal way, by presenting perception as a direct pick-up of information from complex patterns of sensory registrations. On the other side are 'cognitivist' views of perception, represented most recently by the ratiomorphic theories of Rock (1982, 1983, 1985, 1997) and Gregory (1970, 1980), which accommodate the appearance of perceptual givenness under the rubric of propositional inference. I locate perceptual processing, as I located perceptual content, beyond such a dualistic and

oversimplifying schema. Kant's account of perception as a 'receptivity,' as articulated in the *Critique of Pure Reason,* shows us what a third way might look like in this context. Far from reiterating the 'purely passive' model of perception, Kant's concept of receptivity offers a new way of thinking about the passive 'givenness' or 'openness to the world' which intuitively distinguishes perception from thought. When he contrasts the active spontaneity of rational thought with perceptual passivity, Kant is *not* drawing a contrast between a faculty that involves an 'organizing' contribution of the mind and one that does not. Rather, he is contrasting a faculty which *performs* an organizing or interpretative operation *on* something to generate something new; and one in which the organization the mind contributes *precedes* what is organized. More specifically, what makes perception, specifically perceptual constancy, a receptivity is that although the mind supplies a system of organizing principles, the organizing contribution perception makes is not a matter of applying rules to something already apprehended, as when we seek to explain what is perceptually presented by making inferences *from* those presentations. Perceptual capacities are set up in such a way that the organizing structure of a property-space precedes any inputs to perceptions. All the questions are already determined in advance. Perception is a matter of fitting sensory registrations 'in' to that pre-existing framework. With this understanding of perceptual processing, the basis for ascribing a nonconceptual content to perception becomes clear. The nonconceptual, context-dependent nature of perceptual content reflects the limitations inherent in the perceptual process. Perceptual content is context-dependent because the means by which perception acquires information about the external world fundamentally relies on constraints that are particular to the

subject themselves. Perception gives subjects the world from their own, particular point of view.

***

Within the literature on nonconceptual content, it is sometimes claimed that perceptual nonconceptualism admits of two different and logically independent interpretations: 'absolute' versus relative nonconceptualism (Speaks, 2005); or, more commonly, 'content-' versus 'state-' nonconceptualism (Heck, 2000; Byrne, 2005). The distinction traces to a supposed ambiguity in the claim that thoughts have conceptual content; and so a corresponding ambiguity in the denial of a 'conceptual content' for perception. On the first view, thoughts have a conceptual content inasmuch as thoughts are relations to a particular kind of representational content; specifically, a Fregean proposition. In this context, perceptual nonconceptualism is a thesis about the kind of content perceptual experiences have: perception is a relation to a content that is not a structured, Fregean proposition. as claims about a subject's possession of concepts (whether a subject's undergoing experiences with those contents requires them to possess the concepts which characterize those contents). I reject such a distinction. On my understanding of a conceptual content, one could not 'grasp' a constituent of a proposition without 'possessing' the corresponding concept; that is, without possessing the capacities individuated by that structure. Correspondingly, perceptual states, by virtue of having a representational content with a different kind of 'nonconceptual' structure, entail that a subject's capacities will also be

'nonconceptual.' They still possess generality, simply by virtue of being representational. But they lack *conceptual* generality. Perceptual capacities are limited or constrained in their application.[2]

These observations reflect a fundamental bifurcation within the philosophy of mind. For most of the philosophers who play a prominent role in this dissertation—Gareth Evans, John McDowell, Christopher Peacocke, Tim Crane and Tyler Burge—the fundamental entities in psychological explanations are *representational contents*. Philosophers within this tradition are often called ('Neo-') 'Fregean' because they follow Frege in regarding the contents of thoughts and sentences as mind-independent, sharable abstract entities, composed of fine-grained modes of presentation of objects and properties. Within this theoretical framework, representational contents are properties mental states really have, and which contribute to genuinely causal explanations that invoke those states (Peacocke, 1994; Burge, 1995). Yet there are other philosophers for whom the fundamental entities in psychological explanation are not representational contents but entities of a more 'naturalistically respectable' sort, for example, physically realized *representations* (Fodor, 1975) or *biological functions* (Millikan, 1987). The central disagreement here is about fundamentality. What distinguishes the Fregean from the 'Computational Theorist' (one who believes thinking is underwritten by a symbolic 'Language of Thought') is their commitment to the idea that any description of the representational contents of states also

---

[2] The attitude one adopts towards this distinction between 'content-' and 'state-' nonconceptualism is relevant to how one evaluates the adequacy of existing arguments within the nonconceptualism debate. Those who uphold this distinction claim that it undermines the arguments for and against the existence of a nonconceptual content for perception (Speaks, 2005; Crowther, 2006). For a reply to these arguments which reflects the notion of content presupposed by most participants in this debate, see Josefa Toribio's (2007) 'State Versus Content: The Unfair Trial of Perceptual Nonconceptualism.'

provide a complete explanation from the point of view of the kinds of explanations we recognize as distinctively psychological. Many Fregeans have expressed attitudes of neutrality towards the Language of Thought hypothesis (Burge, 2010; Crane, 1992); still others are favorable to this (Peacocke, 1995). Where they agree is that there can be no explanatorily relevant distinctions between properties of representations that do not have a counterpart at the level of representational contents. Claims about the vehicles of representation, and about how they match up with representational contents, pertain to explanation at some other level of description.

Some of the philosophers whose work has had the greatest impact on me are, in one way or another, 'noncanonical.' Amongst them are philosophers who hold nonrepresentational views of the nonconceptual, and whose connections to the nonconceptualism debate have often been overlooked. I discuss this point in chapter one in relation to Travis and Brewer, but there are many others, including Gibson and Merleau-Ponty. I once regarded Travis and Brewer's denial of a 'perceptual content' as, at best, curiously eccentric—and, speaking from my own anecdotal experience, this is still the attitude of most of my colleagues who think about perceptual experience from within the 'Representationalist' tradition ("I just don't get it?" is a refrain I have heard more than once). One has not truly understood their views unless one understands that their stance against representation is fundamentally a stance against conceptual abstraction: what they are opposing (quite correctly, in my view) is a propositional account of perceptual experience. In a rather different vein, I mention also Tyler Burge, whose views about perception most closely resonate with my own convictions. Instead of focusing on

questions about content in a piecemeal way, Burge urges us to look at them in the context

of *how perception (and thought) actually work* in our everyday lives. Indeed, his is the

natural end point for a Fregean view—that differences in the kinds of contents possessed

by perception and thought reflect differences in the ways these capacities work as ways of

connecting us with the world.[3]

Finally, some comments are in order concerning Kant, whom I enlist in this project

as a powerful ally of a nonconceptualist view of perception. Although I am not alone in

reading Kant this way, this certainly has not been the only role assigned to Kant in the

nonconceptualism debate; nor even the most common.[4] If Kant has been widely regarded,

---

[3] On a personal note, I also owe Burge a debt of gratitude for the rigor and unparalleled thoroughness of his writings on perception and representational content. When I first started working on this thesis, I had a hard time finding writers prepared to acknowledge that the structure of perceptual *content*—as opposed to perceptual representations—was even something that it made sense to talk about in relation to nonconceptual content. One by one I rooted out the handful of publications that deal with this question: most notably, Peacocke's chapter on scenario content in his (1992) *A Study of Concepts*, together with papers by Tim Crane (1992, 2009, 2012) and Richard Heck (2000, 2005). But these authors all regard nonconceptual structure as a kind of *unstructured* content. What I could *not* find, even in these writers, and for a long time afterwards, was any clue to substantiate the hunch that has become the central plank of this dissertation: that we might think of the distinction between 'conceptual' and 'nonconceptual' content not as a difference between contents which have a structure and those which lack them; but between structured contents possessing *two different kinds of part-whole structure.* My discovery of Burge's writings on this matter, particularly 'Five Theses on De Re States and Attitudes' (2009) and 'Disjunctivism and Perceptual Psychology' (2005), provided much-needed guidance and encouragement in this regard. They remain the only published characterization of 'nonconceptual classification' of which I am aware.

[4] Indeed, following the rise to prominence of nonconceptual content in the 1990s and his prominence in McDowell's argument for conceptualism, the debate concerning Kant has become a kind of 'sub-sphere' of the nonconceptualism debate. Robert Hanna's Kant is the direct ancestor of Gareth Evans, who defends the existence and meaningfulness of nonconceptual content in a way that "can be directly transferred to the contemporary debate and significantly advance it" (Hanna, 2000; 2001; 2005; 2011; 2016). The disagreement centers in large part on the proper interpretation of Kant's claim of a mutual dependency between concepts and intuitions. Those who read Kant as a conceptualist understand this as claiming that perception, whilst not reducible to thought, requires supplement from concepts to achieve reference to the objective world. Those who read Kant as holding a nonconceptual account of intuition argue that this misreads the context of Kant's claim. According to Hanna, cognition is a technical term with multiple senses: in addition to its broad meaning, as denoting conscious objective representations; a special, narrower sense of 'cognition' as denoting an *objectively valid judgment*. Intuitions and concepts are cognitively complementary and semantically interdependent *for the specific purpose of constituting objectively valid judgments*. But from this it does not follow that there cannot be "empty" concepts or "blind" intuitions outside the special context of objectively valid judgments. (Burge (2010) makes a similar claim). These theorists also point to passages in which Kant

for better or for worse, as a kind of 'founding father' for the conceptualist view of experience, this is due in large part to an infamous passage in which he asserts a mutual dependency between concept and intuition: "Thoughts without content are empty; intuitions without concepts are blind" (A51/B75). In *Mind and World*, McDowell recommends this as a reason why Kant "should still have a central place in our discussion of the way thought bears on reality" (1994: 3). In his introduction to *Body and World*, Dreyfus presents Kant is the very embodiment of the kind of position we should reject:

> The philosophical tradition has generally assumed—or, in the case of Kant, argued persuasively—that there is only one kind of intelligibility, the unified understanding we have of things when we make judgments that objectify our experience by bringing it under concepts. (2001: xv)

Although Dreyfus and Todes disagree with McDowell, holding that Kant's view of perception does justice "neither to the claims of conceptual imagination nor to the claims of perception," they agree with him at least concerning what Kant's view of perception *is*. Even those who do not pursue a conceptualist reading of Kant agree, meanwhile, that Kant's characterization of 'perceptual receptivity' cannot be taken too literally: the concept of 'receptivity' does not, and was never intended to, constitute an characterization of how perception works. In her Introduction to the Hackett edition of *The Critique of Pure Reason*,

---

appears to state quite clearly that intuitions refer to objects in a way that is independent of the semantic and cognitive underpinnings of thought. For example: "*objects can indeed appear to us without having to refer necessarily to functions of understanding*"; the understanding could fail to be objectively valid, yet "*appearances would nonetheless offer objects to our intuition; for intuition in no way requires the functions of thought*" (B122).

Patricia Kitcher presents Kant's distinction between sensibility and understanding, as 'honoring' the common sense idea that there is a difference between perceiving something and thinking about it, without trying to address the difficult—and to this day unanswered—question of 'what exactly the difference between the two faculties comes to.' Kant's distinction between receptivity and spontaneity merely reflects the platitude "that perception involves the sense organs and conception involves concepts"; that "we can have knowledge of the world around us only if we have some faculty for taking in information about that world and some faculty for drawing useful connections between past, present, and future bits of sensory information." Robert Hanna, who reads Kant as a nonconceptualist, refers to the contrast between the spontaneous conceptual functions of the understanding and the receptive perceptual functions of sensibility 'explicatively useful' yet 'misleading' inasmuch as it wrongly implies that Kant's view of perception is "wholly passive or non-generative and non-productive"; rather than being (as Hanna believes) a distinctive *kind* of "lower-level or sensory (receptive) spontaneity" (2001).

I did not set out to engage in Kant scholarship. If I feel compelled to oppose both of these interpretations, it is because I have found the *Critique of Pure Reason* to be not only a reflection of my own nonconceptualist convictions but a continuing source of insight as to the deeper foundations of these. Kant's proposal that the formal structure of intuition gives rise to 'appearances' illuminates, in ways I discuss in chapter two, the non-propositional character of the contents which result from spatial attribution. And Kant's account of the nature of perceptual capacities as a receptivity, which forms the central insight of the account of perceptual processing I propose in chapter four, indicates how the

31

nonconceptual content of perception reflects what is distinctive about how perceptual processing works. Many existing interpretations of Kant overlook what is amongst the most creative aspects of Kant's account of sensory representation. Kant's contribution to our understanding of perception is to show us how perception can involve an organizing of sensory matter without being an operation performed *on* something already given. It is this that distinguishes perceptual processing from the rational, conceptual processes that constitute thinking. Reading Kant in this way allows us to take Kant at his word, when he says that the forms of space and time are *forms of our receptivity*. More importantly still, for present purposes, it allows us to gain from Kant an insight into what a satisfying account of the perception-cognition border might look like.

***

The central thesis of this dissertation is that what it is to perceive the world is to bear a perceptual relation to a representational content that is structured in a 'nonconceptual way.' The four chapters which follow naturally fall into two halves, reflecting two central theoretical issues at stake in the thesis: chapters one and two take up the question, 'What would it mean for perception to have a 'nonconceptual content'?; whilst chapters three and four address the question, 'Why should we believe that perception has a nonconceptual content, understood in this sense?'

Chapter one, 'Reconceiving the Nonconceptualism Debate,' locates nonconceptual content within a broader debate about the nature of perceptual states.[5] The view that perception has a nonconceptual content ('nonconceptual representationalism') is typically opposed to the view that experiences are *sui generis* relations to the same kind of conceptual contents which constitute thoughts ('conceptual representationalism'). The 'reconception' of this disagreement for which I argue consists in the acknowledgement that nonconceptual representationalism and conceptual representationalism do not exhaust the positions theorists have adopted in response to the question of the relationship between perception and conception. There is a third position, 'naive realism,' which holds that experiences differ from thought inasmuch as they are not a matter of representing the world at all; rather, perception is fundamentally a matter of a subject bearing a direct, perceptual relation to mind-independent objects, properties, or events. The naive realist is often regarded as part of a separate debate. The problem with this interpretation, I believe, is that it ignores the motivations of the naive realist, for whom perceptual experience is 'nonconceptual' precisely *because* it is not a matter of representing the world. Drawing on three leading figures from each of these traditions—Gareth Evans, John McDowell and Charles Travis—the chapter draws out the points of agreement and disagreement between these different positions in a way that leads to greater clarity concerning the issues at stake

---

[5] I will often use the term 'perceptual experiences' interchangeably with 'perceptions' or 'perceptual states.' A perceptual experience is a conscious perceptual state. Contrary to what many philosophers have alleged, I take consciousness to be by-and-large irrelevant to the existence and character of representational content in perception. Certainly, I reject the view that equates 'unconscious state' with 'subpersonal state.' Indeed, I am sympathetic to the claim, defended by Rosenthal (2010), that even the qualitative properties of perception are individuated independently of phenomenal properties. The account of nonconceptual content I defend, on which perceptual content has an inherently perspectival or subjective aspect built into it, is naturally in keeping with this claim.

in this debate. The nonconceptual representationalist and the naive realist both agree on the need to sharply distinguish *perception* from *conception*—whilst presenting two opposing views of what it would mean to speak of 'the nonconceptual' in relation to perception. The conceptualist and the naive realist both agree that representation is an essentially conceptual phenomenon—whilst presenting opposing views about what conclusions we should draw from this concerning the objective character of perception. What emerges from this expanded view of the nonconceptualism debate is both an enticement and a challenge for the defender of nonconceptual content. The promise of Nonconceptual Representationalism lies in its potential to reconcile both of the considerations that motivate its rivals: they can acknowledge, with the conceptualist, that the perception and thought are in one fundamental respect relations to *the same kinds of things*; whilst they can also acknowledge, with the naive realist, that perception and thought are in another fundamental respect relations to *different kinds of things*. The challenge is to identify a theory of 'nonconceptual content' that adequately addresses the doubts of those who believe there is no such thing.

Chapter two, 'Classification Without Concepts,' articulates my answer to this challenge: conceptual contents are structured propositions composed of singular and predicate terms; whereas 'nonconceptual content' is a kind of structured representational content which is *non-propositional* and whose constituents are *particulars* and *property spaces*. The argument takes as its foil the 'monist' about classification, according to whom there is only one kind of classification, namely classification through concepts. The chapter centers on a 'substantive' understanding of monism, on which all classification is

conceptual because there is only one form of representational structure that could perform the function of classifying objects as being some way. I argue that this view is false. I also consider a 'deflationary' understanding of monism, which challenges the motivation for regarding this alternative form of classification as 'nonconceptual' content, as opposed to different forms of conceptual content.

The final two chapters mark a turn from purely theoretical considerations towards a practical understanding of nonconceptual content in the explanation of perception. Chapter three, 'Rethinking the Generality Constraint,' takes up the problem of formulating a practical criterion for the existence of nonconceptual content. What kinds of considerations would count in favor of the thesis that a mental state, whether perceptual or of some other kind, has a nonconceptual content? The answer I propose, which draws on Gareth Evans' 'Generality Constraint,' unfolds in relation to my criticism of a more common interpretation of that criterion, which I regard as empirically inadequate. On that interpretation, the justification for the Generality Constraint derives from the alleged systematicity of thought: we ascribe concepts to thinkers to account for the 'empirical datum' that thoughts in fact 'satisfy the Generality Constraint.' But since systematicity is a contingent feature of thought, it is an *open question* whether a creature might have representational states which did not 'satisfy the Generality Constraint.' The main problem with these proposals is that they fail to take seriously the problem of explaining when we are justified in speaking of a subject's *representing* the world at all. The interpretation of the Generality Constraint I offer in place of this presents this as deriving from the connection between *kinds of contents* and *kinds of processes which use and generate representational contents*. We

ascribe conceptual contents to explain a subject's capacity for drawing inferences in which those contents figure as causes. Inferential capacities give rise to a form of *unrestricted generality* in an organism's capacities to successfully discriminate distal features of the environment. Evidence for nonconceptual content would be the presence of a *restricted* form of generality in an organism's capacities to successfully discriminate distal features of the environment, indicating a different kind of underlying process that uses and generates representational contents with a different kind of structure.

Chapter four, 'The Problem of Perceptual Receptivity,' carries these considerations forwards into a concrete proposal about the nature of perceptual processing, which ties this to the theory of nonconceptual content I developed in chapter two. The chapter poses the question: what is it to perceive the world? What separates perception from intellectual thought? The starting point is a compelling intuition, evoked through the image of as a 'receptivity' (Kant, 1996) or 'openness to the world' (McDowell, 1994; Burge, 2010; Brewer, 2006), that what is distinctive to perception as a process of knowledge acquisition, and is that it involves a direct, unmediated relation to the world. The 'problem of perceptual receptivity' emerges when we attempt to characterize this claim in concrete terms. On the one hand, the image of receptivity cannot be taken literally, as 'direct' theories of perceptual processing attempt to do. Perception exhibits too much generality and flexibility to be a purely passive pick up of information from proximal stimulation. Perceptual systems exhibit too much determination and constraint to be regarded as a 'miniature sherlock holmes' drawing on stored 'knowledge' of law-like relations to generate hypotheses about the likely causes of sensory registrations. I argue that Kant's

account of perception as a receptivity, far from being a mere image, offers us a framework within which to reconcile the unavoidable 'givenness' of perception. The process of thought is a spontaneity in which the subject applies their powers of reason *to* something already given. Thinking starts with a perceptual presentation of an object, which the understanding 'scrutinizes to find some rule in them.' Perception is a *receptivity* inasmuch as the organizing structure of space precedes any inputs to perceptions. We have internalized a fixed and unitary 'system' of spatial categories and the laws of euclidean geometry that specify the laws relating those categories to one another. Perception is a matter of fitting sensory registrations 'in' to that pre-existing framework. Building off this idea, I show how Kant's suggestive claims finds a concrete realization in recent predictive processing theories of perception. And I tie these observations about the nature of perceptual processing to a broader picture which reveals their interconnections to properties of representational contents and properties of stimuli in the world.

What emerges is a picture of the nonconceptual content of perception which complements an understanding of perception as a form of 'sensory representation.' Whereas conceptualized thought is a complete abstraction, a grasp of particulars and properties independent of the subject's embodied point of view, perception is essentially embodied, essentially a view of the world from *somewhere*. Perceptual patterning of sensory matter is an imposition of meaning that preserves, in the form of its representational content, the link to the immediacy of sense that is then erased by the propositional structure of conceptual content. Like thought, perception gives us the objective world. Yet inasmuch as it is a form of *sensory* representation, perception gives us

its own world; a world more limited in its reach than, but not lesser than—and certainly not

reducible to—the world of our conceivings or conceptualizations.

1. RECONCEIVING THE NONCONCEPTUALISM DEBATE

1.1 Introduction

Is perception conceptual, or conceptualized? There is an intuitive notion of 'perception' on which this can seem obviously true. We naturally speak of seeing in a way that allows that two people may look at the same tree and see something completely different: one sees a beautiful ancient sycamore tree; whilst the other sees an annoyance that needs to be cut down. Such reports suggest a picture of our perceptual relation to the world as fundamentally *cognitive*, our engagement with objects being always and inescapably colored by the mind's impulse to classify and organize through ever expanding spheres of abstraction. But there is a narrower notion of perception—the one philosophers typically have in mind when they pose this question—on which the answer seems altogether less clear cut. Seeing the tree, in this narrower sense, includes only those of its properties that are *sensed*: its size and vertical orientation, the color and texture of its trunk and leaves, the form of its canopy; in short, those highly stable and seemingly automatic aspects of our awareness of objects that would seem to provide the very soil necessary for our cognitive interpretations to take root. From this point of view, to ask whether perceptual experience is conceptual is to ask not whether the reach of the conceptual is broad but whether it is total. Does the influence of concepts extend all the way down the scale of our sense-based awareness of the world? Or is there some basic layer of perception for which we must acknowledge a role for the 'nonconceptual'?

In contemporary circles, this question has become indelibly shaped by a dialogue between two key figures in late 20th century philosophy of mind. One is Gareth Evans,

whose 1982 monograph, *Varieties of Reference*, is generally credited with introducing the term 'nonconceptual content' into the philosophical discourse. In seeking to understand the role of perception in explaining demonstrative thought, Evans was led to posit a kind of 'nonconceptual informational content' for perception. Thus, corresponding to the distinction between perception and thought we have a distinction between two kinds of representational content:

> The informational states which a subject acquires through perception are non-conceptual, or non-conceptualized. Judgements based upon such states necessarily involve conceptualization: in moving from a perceptual experience to a judgement about the world (usually expressible in some verbal form), one will be exercising basic conceptual skills... The process of conceptualization or judgement takes the subject from his being in one kind of informational state (with a content of a certain kind, namely, non-conceptual content) to his being in another kind of cognitive state (with a content of a different kind, namely, conceptual content). (1982: 227)

John McDowell's John Locke Lectures, published in 1994 as *Mind and World*, are widely recognized as the classic statement of the opposing point of view. Guided by the need to explain the role of perception in justifying our knowledge of the external world, McDowell claims that perceptual experience and thought are different kinds of representational states with the same kind of content in common:

> We should understand… experiential intake—not as a bare getting of an extra-conceptual Given, but as a kind of occurrence or state that already has conceptual content. In experience, one takes in, for instance sees, *that things are thus and so*. That is the sort of thing one can also, for instance, judge. (1994: 10)

Through this disagreement between Evans and McDowell have the battle lines of the 'nonconceptualism debate' been drawn. Philosophers who have followed Evans in defending a nonconceptual content for perception include Christopher Peacocke (1992, 1994, 1998, 2001a, 2001b), Fred Dretske (1982), Tim Crane (1992, 2009, 2012), José Luis Bermúdez (1995, 2003), Richard Heck (2000, 2007), Sean Kelly (2001) and, most recently, Tyler Burge (2010, 2014). Philosophers who have followed McDowell in defending a conceptual content for perception include Bill Brewer (2000) and Alex Byrne (2005). This literature tends to reinforce the idea that the opposition between 'nonconceptual content' and 'conceptual content' is not only central to but *exhaustive of* the views of experience that are at stake in this debate. My aim in this chapter is to challenge that assumption.

In 'Reason's Reach,' Charles Travis starts from the same question that motivated McDowell: what must experience be like so as to justify our beliefs about the world? Travis thinks a view of experience as conceptual all the way down cannot discharge this duty, because conceptual generalities cannot determine what correctly or incorrectly belongs within them. On Travis's view, "to see the meat on the rug is just to be suitably sensitive, or responsive, to it as it then is—to the non-conceptual":

To see *that* the meat is on the rug I must register something else: the instancing by things being as they are of a certain way for things to be, meat being on a rug. I must recognize things being as they are as belonging to a certain range of cases, as what such cases would be. The range in which I thus fit things being as they are is not something visible, nor (present case aside) is what would fit in it. My access to these things is not by sight. ... One does not see (observe) a range requiring one thing or another, any more than one literally sees being meat requiring something for so being.

Travis's endorsement of a nonconceptualist view of experience should not be mistaken for an endorsement of Evans's view of experience, however. In fact, he unequivocally rejects the idea that the nonconceptual is a kind of *representational content*:

Some have spoken of something called 'non-conceptual (representational) content'. On the present idea of the conceptual, that idea makes no sense. Representing as so is essentially conceptual. One represents something as so where there is something it would be for things to be as represented, so what matters to so being, so a range of cases of what would be things so being. With no such range there could be no question   of truth without things being as they are, but then no way for such a question to turn on how they are. (2007: 232-233)

For Travis, then, neither the proposal of Evans nor that of McDowell will do: we must acknowledge a third view of perception, on which experience is a direct relation a perceiving subject bears to the concrete, mind-independent objects. Similar considerations have motivated others, including Bill Brewer (2006) and John Campbell (2002), to adopt the same view. Countering the received wisdom that naive realism is part of a separate debate, I suggest that we should view the naive realist as introducing a new position in the *existing* nonconceptualism debate.[6]

In introducing this third point of view, I am not concerned with promoting it over the positions of Evans and McDowell. As a matter of fact, later chapters will argue against it. My aim is rather to explore how reframing the nonconceptualism debate from a binary disagreement to a tripartite one may deepen our understanding of the complex issues that are at stake. One can reconfigure this triad like a rubix cube, to reveal criss-crossing points of agreement and disagreement. In one configuration, the conceptualist stands opposed by two conceptions of 'the nonconceptual': that of a nonconceptual perceptual *content* and that of a nonconceptual perceptual *relation*. For someone who accepts the need to distinguish perception from conception, what turns on this choice? What is at issue between a view of the nonconceptual as a kind of *content* and a view on which this is objective without yet being a kind of content? In another configuration, the proponent of nonconceptual content stands opposed by two extremes, each of which regard that position as incoherent. What motivates this exclusion of a middle way? What decides in favor of one

---

[6] Whilst I have confined my attention to Travis and Brewer as exemplars of the contemporary debate over representational content, these observations extend to historical or less mainstream strands of the debate over the nonconceptual character of experience. For example, many in the phenomenological tradition, including Merleau-Ponty (1945), Dreyfus (1993; 2001) and Todes (2001) reject a representationalist view of experience *because* they reject a conceptualist view of experience.

or other of these choices? I also draw out two specific consequences of redrawing the battle lines of the nonconceptualism debate so as to include the naive realist within it. The first is that many existing defenses of nonconceptual representationalism have failed to adequately distinguish their position from that of the naive realist. To hold the middle, the nonconceptual representationalist must provide a satisfying account of nonconceptual content which shows how concepts can be subtracted from representational content without thereby leaving us with something that is not a form of *content*. The second consequence is that we have reason to want a notion of the nonconceptual that is distinct from that of the naive realist, because this represents the best chance of reconciling the equally plausible claims made by McDowell and Travis concerning perceptual justification. The challenges raised in this chapter help to set the agenda for my defense of nonconceptual content in the chapters to follow.

1.2 The (Conventional) Nonconceptualism Debate: Evans and McDowell

Thoughts are mental states which have the distinctive property of representing mind-independent objects, properties, and events. Different kinds of thoughts involve different kinds of relations to the subject matter they represent: beliefs aim at those states of affairs that are the case; desires aim at those states of affairs we would like to be realized; intentions aim at those states of affairs we intend to realize; and so forth. The representational content of a thought is associated with three distinctive features: first, it is the fundamental bearer of truth or falsity, and determines truth conditions for the thought it partly constitutes; second, it constitutes the way an individual thinks about a subject

matter; third, it is what is shared in common by different representational states (whether in different subjects or in a single subject at different times) in virtue of which those states are said to be *about the same thing*. The contents to which subjects are related in thought are 'conceptual contents.' I shall assume that a conceptual content is a Fregean proposition: a complex intentional object having a 'sentence-like' structure; and whose constituents are fine-grained *modes of presentation* of objects and properties. Correspondingly, I take a concept to be a constituent of a Fregean proposition.[7] Concepts possess three semantic properties which complement the semantic properties of the representational contents they compose. First, a concept helps constitute a truth condition, in the sense that the overall proposition is true or false according to facts about the referents of its constituents, together with their manner of combination. Second, a concept constitutes a (fine-grained) way in which a subject thinks about objects and properties. For example, the concepts *coriander* and *cilantro* present the same referent in different ways; one who thinks of coriander (cilantro) as *coriander* thinks of it under a different guise or in a different way from one who thinks of it as *cilantro*. Third, a concept is what is shared in common by different thoughts which are (partly) about the same thing. Concepts thus provide the basis for comparisons of cognitive states within a given agent and across agents. Many philosophers within the Fregean tradition associate concepts with a fourth property: they

---

[7] The Fregean view of concepts is controversial. Within the 'computational' tradition concepts are identified not with constituents of *propositions*, but with the constituents of structured, symbolic *representations* (i.e. the vehicles of content). Conceptual contents are thought of derivatively, as the contents of sentence-sized representations (Fodor, 1975; Carey, 2009; Gawker, 2011; Pylyshyn, 1973, 2007). For those in the computational tradition, conceptual contents are typically understood to be Russellian propositions, whose constituents are objects and properties themselves. Although the disagreements between these two frameworks will figure in subsequent chapters, I shall ignore it for the purposes of the present chapter. My aim is to understand the mainstream nonconceptualism debate; and amongst the participants in that debate (both defenders and critics alike), the Fregean view is accepted by virtually all.

mark psychological capacities on the part of the subject who thinks with them. On a strong version of this view: if a subject thinks that *a is F*, that subject also has the capacity to think the thought that a is G (for an arbitrary predicate concept, G, which that subject possesses); and likewise, they have the capacity to think the thought that b is F (for an arbitrary singular concept, *b*, which that subject possesses) (Evans, 1982; Peacocke, 1992).[8]

The nonconceptualism debate starts from the assumption that perception is a *sui generis* kind of mental state: like a belief, a perceptual experience is directed on the environment; but perceptual experiences are not reducible to beliefs or other kinds of thoughts.[9] The question that is centrally at issue in the nonconceptualism debate is whether the differences between experiences and beliefs entail differences in how experiences present the world as being. Shall we think of the differences between perception and thought as purely a function of their being different kinds of (conceptual) states? Or must we acknowledge, in conjunction with the *sui generis* character of perception, a *sui generis* kind of perceptual content? Conventionally, discussions of the nonconceptualism debate describe two possible positions one can take with respect to these questions. The nonconceptual representationalist answers in the affirmative: experiences are relations between subjects and nonconceptual contents; contents of a different kind from those that characterize our thoughts about the world. The conceptual

---

[8] Various aspects of this view about concepts are discussed more fully in later chapters. See chapter two for a more complete discussion of concepts and 'sentence-like' structure. See chapter three for a detailed discussion of the Generality Constraint.

[9] The position that perception is a sui generis representational state stands in contrast to 'belief theories' of perception, as defended by Armstrong (1968) and more recently, Gluer (2009). On these views, perception is the acquisition of a belief, or a disposition to acquire beliefs (an acquisition being a conscious event, as perceiving is; rather than a state or condition, as belief is.)

representationalist, meanwhile, answers in the negative: experiences are different from

beliefs, but the contents of the two kinds of representations are of the same kind.

The original statement of nonconceptual representationalism within contemporary

philosophy of perception can be traced to Gareth Evans's 1982 monograph, *Varieties of*

Reference. Evans introduces the concept of perception as part of the 'informational system':

> When a person perceives something, he receives (or, better, gathers) information
>
> about the world [... This locates perception] in a system—the informational
>
> system—which constitutes the substratum of our cognitive lives.
>
> A traditional epistemologist would have recast these platitudes in terms of
>
> the concepts of sensation and belief. (...) (T)he subject would have been regarded as
>
> receiving data, intrinsically without objective content, into which he was supposed
>
> to read the appropriate objective significance by means of an (extremely shaky'
>
> inference. [...] (I)t is now widely realized that the traditional conception gets things
>
> impossibly the wrong way round. The only events that can conceivably be regarded
>
> as data for a conscious, reasoning subject are seemings—events, that is, already
>
> imbued with (apparent) objective significance, and with a necessary, though
>
> resistible, propensity to influence our actions. (1982: 122-123)

As this passage illustrates, Evans's interest in perception derives primarily from his larger

interest in explaining demonstrative reference in thought. For Evans, the technical notion

of an 'informational state' is rooted in what he calls 'information-based thought,' a special

subset of thoughts.[10] In information-based thoughts, 'the overriding point or purpose of the

subject's thinking is to be thinking of the object from which the information derives.' One of

the moves being made in this passage is to establish perception as a state of a certain kind,

namely one belonging (together with memory and communication) to the 'informational

system.' The states in the 'information system' are those which Evans identifies as crucial

to explaining information-based thought. For example, a person may be thinking of an

object which he can perceive—for example, a black and white cat sleeping on a mat.

Assuming that this person has no information about the cat other than that which he is

acquiring by current perception, the conception governing his thought will be determined

simply by the content of his perception, in the sense that it is with reference to the

perceptual content that the contents of his thoughts will be determined. There is a special

case of information-based thought for which perceptual information plays a particularly

crucial role, namely demonstrative thought. A demonstrative thought, 'the mother and

father of all information-based thoughts,' is one involving a *continuing* information link; the

subject's thinking is governed by a controlling conception of an object he can pick out by

sight or hearing or touch, or otherwise sensibly discriminate ("If the question were raised,

'How do you know there is such a thing as the thing you take yourself to be thinking about?'

he would answer can 'see' (or 'hear,' or 'taste,' or 'feel') that there is" (ibid., 146)). As Evans

---

[10] Evans identifies information-based thoughts thoughts as having certain distinguishing features; most notably, they are "liable to a quite distinctive failing: that of being ill-grounded. ... this fact about information-based thoughts has considerable consequences for the theory of reference; for there are many uses of referring expressions for the understanding of which an information-based thought is required" (ibid., 134-135). He regards this subset of thoughts as playing a critical role in explaining a range of cognitive and linguistic phenomena, including demonstrative and recognition-based identification (as opposed to identification by description), sentences containing ordinary demonstratives, referring to an object in the shared perceptual environment; sentences containing 'past-tense demonstratives', referring to an object earlier observed and now remembered; and sentences containing 'testimony demonstratives', which advert to information presumed to be in common possession through testimony.

comments in relation to spatial demonstrative thought, "(w)here there is no possibility of action and perception, 'here'—thoughts cannot get a grip" (ibid., 153).

The foregoing commentary establishes a general theme in Evans's approach to perception. *Varieties of Reference* is, first and foremost, a book about thought; and Evans's interest in perception derives from what he regards as its crucial role in explaining thoughts of certain kinds. As he frames his guiding question: "we need an explanation of exactly how it is that perceiving something makes a thought of a certain kind possible." One part of Evans's answer to this question is that informational states in general, and hence perception in particular, are not reducible to beliefs, judgements, or other forms of thought. The claim that perceptions are "seemings," in the foregoing passage, is a first step in this direction. The point here is that perceptions are a kind of representation that is *essentially sensory*: they are representations in which sensation and 'objective significance' are intertwined (and not, as on some older theories of perception, two separate steps in an inferential process).[11] Evans invokes two further considerations which support a sharp distinction between perception and thought. One is that perception is 'belief-independent' (or 'cognitively impenetrable'), in the sense that what we perceive is not affected by beliefs or other thoughts. In support of this claim, he gives the example of the Muller-Lyer illusion, in which we continue to perceive one line as being longer than another even when "we are quite sure that it is not." The second consideration is prompted by reflection on 'dispositional' accounts of perception, which identify perceptual states with dispositions to

---

[11] A famous example of the inferential view of perception would be Berkeley's constructivist view of perceptual constancy (1709, 1733). More recently, the inferentialist view has been defended by a number of cognitive scientists, including Irvin Rock (1982, 1983) and Richard Gregory (1974).

form beliefs. Evans's pronouncement is that "this gets things the wrong way around": the dispositional account tries to subsume perceptual states under a more sophisticated cognitive state that is implicated in judgement and reasoning; whereas the natural view is that perception is a more primitive kind of state which "we share with animals" and from which capacities for rational thinking presumably evolved (ibid., 124). Notice that these claims are not meant to be specific to perception; Evans regards them as applicable to the 'informational system' at large.

This pair of claims—first, that perception is an informational state and second, that perception (*qua* informational state) is distinct from thought—provide the framework for Evans's claim that perception has a 'nonconceptual informational content.' At the core of our idea of perception is, as he puts it, the idea "of an information-link between subject and object, which provides the subject with (non-conceptual) information about the states and doings of the object over a period of time" (ibid., 130). Evans thinks that the differences between informational states and beliefs in and of themselves support a difference in the natures of their representational contents: informational states have a distinctive kind of 'informational content' involving 'a certain bit of information being of, or perhaps from, an object'; as distinguished from 'the notion of belief about an object, or thinking about an object'. Moreover, Evans holds that because of their special role in explaining demonstrative thoughts (specifically, spatial demonstrative thought such as 'I am here,' 'it's F over there') perceptual information must take a special form. Specifically, perceptual information must be embedded within an egocentric spatial frame of reference in which the subject is positioned,

at its center (at its point of origin), with its co-ordinates given by the concepts 'up'

and 'down', 'left' and 'right', and 'in front' and `behind'.... Egocentric spatial terms

are the terms in which the content of our spatial experiences would be formulated,

and those in which our immediate behavioural plans would be expressed. (1982:

154)

A subject does not merely hear a sound; rather, a sound is always heard as coming from

such-and-such a position in space, *from* a location relative to the *here* that is the origin of

the subject's egocentric reference frame. Evans identifies this spatial element of perceptual

content as a central case of nonconceptual content in perception. A final reason Evans

mentions for recognizing nonconceptual content pertains to the fine-grained character of

perceptual experience. Perceptual experience is fine-grained inasmuch as an experience of

(say) a wooden joist presents its shape and size, texture and orientation in a way that

specifies highly determinate values along each of these dimensions. According to Evans,

no account of what it is to be in a non-conceptual informational state can be given in

terms of dispositions to exercise concepts unless those concepts are assumed to be

endlessly fine-grained; and does this make sense? Do we really understand the

proposal that we have as many colour concepts as there are shades of colour that we

can sensibly discriminate? (ibid., 229)

It is a requirement on an account of perceptual content that it capture the way the world seems to us in experience. What Evans is suggesting here is that this requirement cannot be met using only those contents which can be built up by referring to concepts of the properties which the experience represents. Even the most fine-grained color concepts we possess, the concept of *alizarin crimson* for instance, would not come close to describing the contents of our experiences of colors in a way that does justice to their fine-grained character.[12]

The classic articulation of the conceptual representationalist position is John McDowell's *Mind and World* (1994). McDowell's starting point is a Kantian problem: how do concepts, hence conceptual thoughts, make contact with the world? The problem gets its bite through a consideration of two equally problematic proposals about the relationship between thought and perception. One is the doxastic view, which reduces perception to a judgement or belief about sensation. The problem with this view, according to McDowell, is that since judgements are part of the realm of rational freedom, there is a risk that thinking becomes "a play of concepts without any connection with intuitions, that is, bits of experiential intake" (ibid., 4). As Kant points out in the *Critique of Pure Reason*, "I can think whatever I please, provided only that I do not contradict myself; that is, provided my concept is a possible thought" (A569n/B624n). It is supposed to be their connection with experiential intake that supplies the content, the substance, that thoughts would otherwise lack: There must be external constraint if our activity in empirical thought and judgement is to be recognizable as bearing on reality at all. The failures of the doxastic view give rise

---

[12] It is worth noting that, although this is arguably the most famous of Evans's claims, because of the subsequent debate it has generated, it is mentioned only in passing, and is not part of Evans's central commitments.

to a second, equally problematic suggestion: that perception is a 'given,' a kind of raw data which somehow pertains to the world without yet being a representation of that world. When McDowell labels the Given as a 'myth,' he charges it with being 'useless for the purpose it was introduced.' The myth is that we could resolve the problem by locating experience outside the realm of what is intelligible in thought. In so doing, we locate it outside of the realm of what can meaningfully be said to constitute *reasons* for our thoughts about the world: 'bits of experiential intake' can be at best be causes of, and not reasons for, judging things to be one way, rather than another. For this reason, McDowell rejects any view of experience that attempts to "make use of the notion of an interface between mind (which inhabits the space of concepts) and world, where the world presents the mind with non-conceptual items for it to work into conceptual shape" (1994a: 205).

This, then, is the dilemma which lays the groundwork for a conceptual view of experience: we cannot accept the view of experience which makes this a 'given'; and nor can we fall back on the view which prompted it. The problem is resolved, McDowell claims, by eliminating the dualism between 'scheme' and 'Given.' Experience can play the role of a 'tribunal' for belief precisely because in experience there is *no distance* between the impressions from the world and the conceptual contents that are taken up into judgement:

> The conceptual contents that sit closest to the impact of external reality on one's sensibility are not already, *qua* conceptual, some distance away from that impact. They are not the results of a first step within the space of reasons, a step that would be re-traced by the last step in laying out justifications, as that activity is conceived

within the dualism of scheme and Given. This supposed first step would be a move

from an impression, conceived as the bare reception of a bit of the Given, to a

judgement justified by the impression. But it is not like that: the conceptual contents

that are most basic in this sense are already possessed by impressions themselves,

impingements by the world on our sensibility. (ibid., 12)

On a conceptualist view of experience, experience and judgement are brought 'in

alignment' with one another through their shared content, whilst the essentially sensory

aspect to perceptual representations allow us to clearly distinguish perception from

judgement. Unlike judgement, perception is a receptivity to the world, that is, an 'openness

to manifest facts, facts that obtain anyway and impress themselves on one's sensibility"

(ibid., 29). Moreover, the singular components of genuinely perceptual contents are object-

dependent demonstrative senses. Yet the conceptual contents which partly compose

experiences provide the semantic structure or logical form that can make rational the

transition from experience to belief. The capacities that are drawn on in experience are

conceptual go hand in hand with 'responsiveness to rational relations' which link the

contents of judgements of experience with other judgeable contents; "these linkages give

concepts their place as elements in possible views of the world" (ibid., 12).

McDowell's opposition to nonconceptual representationalism is that it is a 'dressed

up' version of the myth of the given (ibid., 53). Evans's view can seem to escape the worries

that dog the notion of a given, by combining the denial of conceptual content with an

insistence on some kind of perceptual content. This makes it seem as though what we have

in experience is something capable of being answerable to the world and upon which judgements could be 'based.' In placing perceptual experience 'across a boundary that encloses conceptuality,' however, Evans places experience in a position that makes this impossible: for what is not conceptual cannot be normative; and what is outside the normative sphere is what is outside of the sphere of what could meaningfully be called a 'content' at all. In short, nonconceptual content is a 'fraudulent labelling,' where the term 'content' refers to something fundamentally unlike what we mean when we apply that term in relation to thoughts. Supplementing this central line of argument concerning epistemology, McDowell also challenges the adequacy of Evans's arguments for nonconceptual content. For example, Evans's claims about the fine-grained character of experience overlook an alternative explanation for these facts: any perceived colour shade, no matter how fine-grained, can be captured in terms of the demonstrative concept *that shade* (ibid., 53).[13] Thus, Evans draws an *invalid* inference from differences in representational states to differences in representational contents: whilst nonconceptual content is one explanation for the differences in question, it is not the case that these features *can only* be explained by adverting to a nonconceptual content. The general point

---

[13]The same points have been made against nonconceptualist arguments from the apparent overlap between human perception and those of nonconceptual animals and infants, which overlook explanations in terms of subpersonal processing. Jerry Fodor criticises the nonconceptualist for illegitimately assuming a link between concepts and propositional thoughts: if a creature lacks propositional thoughts, they also lack concepts. This fails to consider, as an alternative explanation for the differences between perception and thought, that perception is 'modular' or 'informationally encapsulated': encapsulated contents are those that can't interact with any other 'higher' cognitive processes; hence the concepts deployed in modular subsystems can't be used in inferences made in central cognition. Directing his criticisms against Tyler Burge, a prominent proponent of Nonconceptual Representationalism, he charges Burge with 'begging the question' against modularity theory: "Burge needs an argument why *only* constituents of propositional thoughts can be concepts; why shouldn't concepts also, and independently, have a role in the perceptual integration of sensory registrations? This seems to me like the paradigm of a question that can't be settled from the armchair" (2015: 212).

behind these criticisms is that nonconceptual content, apart from being incoherent, is simply unmotivated.[14]

The dialogue between Evans and McDowell is responsible for shaping, to a powerful degree, the subsequent debate. Evans's influence is evidenced in the theoretical concerns which motivate a nonconceptual content to perception: the role of nonconceptual content in grounding demonstrative reference in thought; his idea that we need a less demanding or sophisticated notion of nonconceptual content to capture the representational capacities of creatures incapable of rational thought and language (Peacocke, 2001; Burge, 2010); the 'belief-independence' of perception (Crane, 1992; 1988); and (more contentiously) the fine-grained character of experience (Peacocke, 1986; Kelly, 2001). Evans's account of nonconceptual content as an egocentric frame of reference is also echoed in Christopher Peacocke's theory of 'scenario content,' the most influential and complete account of nonconceptual content. On the other side of this debate, conceptualists have followed McDowell in questioning the coherence of nonconceptual content. Alex Byrne, in a variation on McDowell's appeal to experiences as reasons, comments:

> McDowell disparages nonconceptualism as another version of the Myth of the Given, and the comparison is particularly apt. The traditional Given is ineffable, a feature

---

[14] From the point of view of a contemporary observer of the nonconceptualism debate, Evans's discussion of perceptual nonconceptual content can feel dissatisfying, if not outright irresponsible. Evans effectively sneaks the thesis in 'under the radar' through two other suggestive, though inconclusive, observations: that perceptions stand in a certain relation to some kinds of thoughts; and that perceptions are different kinds of states from thoughts. But it is not immediately obvious that either of these facts cannot be explained in terms of a view of perceptual experience as a *sui generis* conceptual state. Whilst I think there is something fair in these criticisms, we must also note that since there was no 'nonconceptualism debate' at the time Evans was writing his claims didn't carry quite the same burden of proof. Moreover, in his brief discussion Evans hits virtually all the issues that have become key considerations in the subsequent debate.

shared by nonconceptual content. The nonconceptual content of experience is not thinkable—and it cannot be whistled either. Reflecting on one's experience, one might have some inchoate suspicion that there is something special about its content, and often this seems to motivate nonconceptualism. Yet any such motivation is doubtfully coherent. (2005: 23)

McDowell's criticism of the adequacy of Evans's arguments for nonconceptual content has also become a recurring theme (Speaks, 2005; Crowther, 2006). Some have martialed these observations into an argument for conceptualism, by supplementing them with considerations of explanatory parsimony: other things being equal, conceptualism is the default view because it carries fewer ontological commitments; and since nonconceptualists have not shown that there are phenomena which absolutely cannot be explained save by appeal to their view, we have no good reason to abandon conceptualism (Byrne, 2005; Robbins, 2002; Duhau, 2009). Notice that this approach in some ways distorts McDowell's original intentions, however. The idea which animates McDowell's conceptualism is not simply that we should assume a conceptual content for perception because no compelling reasons have been offered for why we should not. Rather, his motivations are rooted in a commitment to a robust view about the role of representational content in explaining the nature of mind, and a positive vision of what experience must be like to fit within this picture.

1.3 The Reconceived Nonconceptualism Debate: Travis

Let us return to the question that is centrally at issue in the nonconceptualism debate: 'Does perception have a conceptual content?' We have considered two opposing answers to this question: the conceptual representationalist regards perception as conceptual all the way down; the nonconceptual representationalist regards perception as a relation to a nonconceptual content. I should like to propose that we acknowledge, as part of this debate, a rival conception of the 'nonconceptual' in relation to perception. On this view, by virtue of involving a *non-representational* or *direct* relation a subject bears to mind-independent objects and properties. I shall refer to this view as naive realism. According to the naive realist, it is *precisely because* the representational model is suitable for thought that it is unsuited to experience. What distinguishes perception from thought is that 'it is the actual physical objects before her which are subjectively presented in a person's perception, rather than any of the equally truth-conducive possible surrogates.' (Brewer, 2006: 174).

One observation naive realists have made, in opposing a representationalist view of experience, is to point to the inadequacy of arguments for that view. Consider the following question posed by Susanna Siegel:

Is it already part of your visual experience that John Malkovich is walking by, carrying a dog? Or do you just visually experience an array of colored shapes bouncing slightly at regular intervals, and subsequently judge that it is John Malkovich carrying a dog? More generally, we can ask: do you just visually

experience arrays of colored shapes, variously illuminated, and sometimes moving? Or does visual experience involve more complex features, such as personal identity, causation, and kinds such as bicycle, keys, and cars? … I will argue … [that] the contents of visual experience are richly complex, and so are not limited to color, shape, and other properties standardly taken to be represented in visual experience. (2010: 3-4)

Travis accuses Seigel of tacitly assuming an equivalency between a theory-neutral notion of the 'objective character' of perception and a theory-laden notion of perceptual content:

Either perceptual experience has representational content, and moreover rich content in her sense; or what we experience is merely shapes, colours, movements (etc.). Whether perceptual experience has representational content or not, and if so, whether rich or not, is, from the start, a matter of what its objects are—whether, e.g., John Malkovich or not. Or thus according to Siegel. There is no room from her perspective for a view on which what one sees if John Malkovich walks by holding a dog is: John Malkovich walking by holding a dog (unless one's attention was diverted at the crucial moment, or all went blurry, or there was excess of glare, or etc.); but representational content belongs only to the attitudes one forms in response to this. (2012: 2)

If experience is 'directed on' or 'of' mind-independent objects and/or properties (so the assumption goes) then there can be no objection to the view that experience is fundamentally a matter of representing the environment as being a certain way—because the two claims are just different ways of saying the same thing.[15] But this assumption ignores the fact that the 'content view,' on which 'ofness' or 'aboutness' is specified by an abstract object which functions to match the world, is a *theoretical model* for how to capture the objective character of mental states. Moreover, it is only one of several possible such theoretical models. For one might also claim, as the naive realist does, that perception consists of a *direct* relation between world and mind, unmediated by an abstract representational content. To reject the claim that perceptual experience has a content is consistent with acknowledging the fundamentally objective status of perceptual experience.

Because their criticism targets an assumption that is made by both parties to the nonconceptualism debate, one might suppose that the naive realist is part of a separate debate. If the naive realist's objection is to an assumption that is common to conceptual representationalist and the nonconceptual representationalist, the question of whether that assumption is correct is orthogonal to the question at issue between them. But this narrative overlooks the deeper motivations the naive realist has for opposing a representational content for perception. These are inseparable from questions about the kind of content perception has. In order to bring out this point, let us consider Travis's

---

[15] In fairness to Seigel, and as Travis himself concedes, she is one of the few who actually tries to argue for a representational content for experience. But one does not have to look far to find examples of philosophers who do not.

argument in 'Reason's Reach.'[16] We may take as our starting point Travis's defense of a

metaphysical distinction which he calls 'Frege's Line,' on which his argument against the

Representationalist crucially turns. He introduces this distinction by means of the following

quotation:

> But don't we see that the sun has set? And don't we thus also see that this is true?
>
> That the sun has set is no object which sends out rays that reach my eyes, no visible
>
> thing as the sun itself is. That the sun has set is recognized on the basis of sensory
>
> impressions. For all that, being true is not a perceptually observable property.
>
> (2007: 229)

On the basis of this passage, Travis's first characterization of Frege's Line marks a

distinction between two kinds of objects: on the left lie *sensible* objects (the sun, and things

which, like it, reflect or emit light into one's eyes); whilst on the right lie *intentional* objects

or representational contents (*that the sun has set*, and things like it which are evaluable as

true or false). The point of calling it a line is that the distinction is exclusive. That something

is a sensible object excludes its being a representational content, since a sensible object—

like a particular piece of meat or the particular computer I see in front of me, is not the sort

of thing that can be evaluated as true or false: "A thing represents only insofar as it is to be

taken as representing in a particular way (as opposed to others in which, for all its

---

[16] It may be noted, however, that the general commitments Travis articulates in this paper, and which motivate him to reject a representational content for experience, are not idiosyncratic to Travis. In his 2006 paper, 'Perception and Content,' Bill Brewer motivates a naive realist view of experience on the basis of similar claims. More-or-less all of what I shall say about Travis equally applies to Brewer. (See also: p.63, fn16)

perceivable features, it might). … It is in this sense that what is perceivable is excluded from 'the domain of things for which truth can come into question at all.'" Inversely, if something is a representational content, it is not the sort of thing that could be sensed, because sensible objects are objects that can *causally affect the senses*; propositions, being abstract objects without a spatiotemporal location, are not the sorts of things that could casually impact sense organs.

Following this, Travis offers a second characterization of Frege's Line:

To the right of Frege's line is the conceptual. What is there to the left? What instances (first-order) conceptual generalities. Such as that piece of meat. (ibid., 232)

For Travis, a concept is the semantic analog of a property: just as a property (being red meat) is an abstract object which can be instantiated by a range of different particulars, so a concept (the concept of being red meat) is an abstract object which can 'fit' (be asserted of or predicated of) *a range* of qualitatively distinct particulars. As he puts it, "the key feature of the conceptual… is that for anything conceptual there is a… range that is the range of cases, or circumstances, which would be ones of something instancing that generality (or, again, a range of things not instancing it)" (ibid., 231). Travis designates the right the concepts and objects which 'instance' conceptual generalities (like *that piece of meat*). This distinction is again exclusive: if something is capable of *instantiating* conceptual generalities then it is not the sort of thing that could be a *instantiated by* things; and if

something is a concept then it is not the kind of object that could instantiate concepts. On Travis's usage then, 'Frege's Line' does double duty by standing for *two* opposing pairs of objects: sensible objects versus representational contents on the one hand; and objects which 'instantiate' conceptual generalities versus concepts on the other. Why does he identify concepts and contents, which are ostensibly distinct? Because the generality is in a sense the same kind of generality in both cases: at the level of wholes and at the level of parts: "A thought always contains something reaching out beyond the particular case, by which this is presented to consciousness as falling under something general." The conceptual and propositions belong to the right of Frege's line because they *belong together*. And likewise, what instances (first-order) conceptual generalities lies on the same side as sensible objects because these things belong together: they are truth makers, the things which concepts and propositions aim to fit or match, and relative to which their fitting or matching may be judged a success or not. Thus, the identification is no slip of the pen. Rather, it marks the fact that representational contents and concepts both reside on the right side of Frege's line, that a representation as so is *essentially conceptual*.

If we accept Frege's Line, then the theorist of perception is presented with a choice between two kinds of entities to which subjects are related in perception. The first is that the objects of perception are those entities which fall to the left of Frege's line: intentional objects, or (more specifically) propositions, or (more specifically still) conceptually structured propositions. The other is that the objects of experience are those entities which fall to the right of Frege's line; that is, concrete, physical (sensible) objects. There are no other options. The first choice corresponds to McDowell's view of experience. The second

choice corresponds to Travis's view of experience. And Travis's argument for this

conclusion is in effect to turn McDowell's desiderata against him: the only possible way for

experience to bear on conceptual generalities, to count for or against their correct or

proper application, is if experience is a relation to the *nonconceptual*. Hence, "there

precisely must be rational relations between the conceptual (what satisfies the condition)

and something else if we are to make sense of experience bearing on what one is to

think".[17,18]

Notice, however, that in disagreeing with McDowell in this way, Travis is not

thereby agreeing with Evans. The nonconceptual does not mean for Travis what it means

---

[17] Travis offers a further reason why we ought not accept a view of experience which locates the objects of experience to the right of Frege's line: One is that he thinks there is something inherently incoherent about claiming that representational contents are the sort of things that can be seen. For, given the observations adduced earlier, representational contents are not *sensible* objects: "The range in which I thus fit things being as they are is not something visible, nor (present case aside) is what would fit in it…we [must disallow] *that the meat is on the rug* as an object of visual awareness" (Travis, 2007: 232).

[18] Bill Brewer's 'Experience and Content' argues against representationalism based on similar theoretical commitments: the main error of the Representationalist is "a mistaken importation of selective intellectual abstraction, or categorization, into the account of perception." The Representationalist is committed to the idea that your perceptual experience has specific truth conditions, which go beyond anything fixed uniquely by the actual nature of the [perceptual object]. According to the representationalist, then, perception,

> offers a determinate specification of the general ways such constituents are represented as being in experience: ways which other such constituents, qualitatively distinct from those actually perceived by any arbitrary extent within the given specified ranges, might equally correctly—that is, truly—be represented as being. Any and all such possible alternatives are entirely on a par in this respect with the object supposedly perceived, so far as (CV) is concerned. Thus, perceptual experience trades direct openness to the elements of physical reality themselves, for some intellectual act of classification or categorization.

Brewer regards this commitment of Representationalism to conceptual generality in experience as unacceptable: "The selective categorization of particular constituents of physical reality enters the picture of a person's relation with the world around her only when questions of their various similarities with, and differences from, other such things somehow become salient in her thought about them, rather than constituting an essential part of their subjective presentation to her in perception"; perception itself "constitutes the fundamental ground for the very possibility of" such abstract classification of experience in conceptual thought. (Brewer, 2006: 174) Once this point is recognized, it allows us to understand how it could be that Brewer changed his position from a conceptualist view of experience in Sense and Content (2000) to a nonconceptualist account (2006).

for Evans. For Evans, experience has a 'nonconceptual content'. But, if we follow Frege's Line, that is impossible, because propositions and concepts belong together. The claim that perception is nonconceptual—that experience does not reside on the right hand side of Frege's line—can only mean that it resides to the left, in the realm of concrete objects and properties that are the *truth-makers* for propositions Experience is nonconceptual in the sense that one is experiencing (witnessing, seeing) things as they are, with no commitment to their being *one way or another.* For Travis, experience "confronts us with what is there, so that, by attending, noting, recognizing, and otherwise exercising what capacities we have, we may . . . make out what is there for what it is—or, again, fail to . . . in perception things are not presented, or represented, to us as being thus and so. They are just presented to us, full stop" (Travis, 2004: 65).

We are now in a position to appreciate why the disagreement the naive realist has with the representationalist is not independent of the nonconceptualism debate. We have seen that the rejection of a representational content for experience is predicated on a certain understanding of what representational contents are, what theoretical commitments accompany the ascription of a content to a mental state. Representational content is *essentially conceptual* in the sense that to speak of a representational content is to speak of something that by its nature can 'apply to' or 'fit' a range of distinct entities, by virtue of which the content is evaluable as correct or incorrect. But perceptual experience cannot be general in this way, which means that perceptual experience cannot have a representational content. For one may refute this line of argument in one of two ways: either they are wrong to claim that the conceptual contents of experience pose a problem;

or they are wrong to claim that the contents of experience must be conceptual. To refute the argument, the representationalist must argue for one or other of these claims. But to argue for the first claim is to side with the conceptual representationalist, whilst to argue for the second claim is to side with the nonconceptual representationalist. Hence, the representationalist cannot engage with the naive realist without taking a stand on the question whether the contents of experience are conceptual or nonconceptual. Worth noting, in conjunction with this point, is the fact that the naive realist's disagreement with the conceptual representationalist is on a different basis from the disagreement they have with the nonconceptual representationalist. They reject a conceptual content for experience because 'the generality of the conceptual' is fundamentally unsuitable for capturing the nature of perceptual experience. They reject a 'nonconceptual content' for experience because they regard the notion of a nonconceptual content as *incoherent*. To put it another way, in opposing representationalism, the naive realist has two opponents. Their *direct* opponent is specifically the conceptual representationalist. But in characterizing the nonconceptual in terms of the constituents of the physical world, rather than in terms of a kind of representational content, they position themselves in *indirect* opposition to someone like Evans—whose rival account of the nonconceptual they regard as inconsistent with the nature of representational content. To counter the naive realist, either one must reject their claim that experience cannot be a conceptual content, or one must reject their claim that the nonconceptual cannot be a form of content. Which is just another way of getting at the point that one cannot engage with the naive realist without taking a stand in the (conventional) nonconceptualism debate.

1.4 Two Perspectives on The Nonconceptual

Recognizing the naive realist as part of the nonconceptualism debate has implications for how we evaluate the dialectic between conceptual and nonconceptual representationalism.[19] This section will explore the disagreement between Evans and Travis concerning the nature of the nonconceptual against their background agreement that an account of perception must be able to accommodate the particularity of perception. In the next section, I consider the disagreement between McDowell and Travis in the context of their background agreement that merely acknowledging the naive realist as part of this discussion provides a prima facie motivation for locating a viable notion of 'nonconceptual content.'

We have seen that Travis's objections to a representational content are indebted to Frege. But in fact, I think a consideration of Frege's views about perception can be useful to understanding the position of Evans as well as Travis, and the relationship between them. In 'The Thought,' Frege writes,

---

[19]The conclusion also has implications for the existing debate between naive realists and representationalists, since it suggests that many of the arguments that have been given in favor of representationalism are inconclusive.In response to complaints from naive realists that "few … actually bother to argue that perceptual experience does have representational content" (Travis, 2012: 1), a number of philosophers have given arguments in favor of representationalism (Byrne, 2009; Seigel, 2010; Schellenberg, 2011; Pautz, 2009). A common feature of arguments in defense of representationalism is that they ignore the question of what the nature of perceptual content is, operating instead with a notion of representational content that is neutral between conceptual and nonconceptual conceptions thereof. They thereby ignore the main objection to a representationalist view of experience that is motivating the naive realist: namely its commitment to a conceptual content for perception. they fail to address the central motivation for naive realism. They are therefore talking past their opponents rather than to them. The representationalist cannot advance this debate without reflecting more carefully on the nature of perceptual content. Travis's review of *The Contents of Visual Experience* consists largely in a restatement of the points made in 'Reason's Reach'—ignored by Seigel in her discussion of Travis. Since a discussion of this point would take us beyond our present concerns, I shall not pursue it here.

> To have visual impressions is not to see things... Having visual impressions is certainly necessary for, seeing things but not sufficient. What must still be added is non-sensible. And yet this is just what opens up the outer world for us; for without this non-sensible something everyone would remain shut up in his inner world. So since the answer lies in the non-sensible, perhaps something non-sensible could also lead us out of the inner world- and enable us to grasp thoughts where no sense-impressions were involved. (1918/1956: 305)

This passage is offered as an answer to a certain problem Frege poses for himself: "How does it happen that I see the tree just there where I do see it?" That question arises, in turn, through a distinction between the 'outer world' (or objective realm) and 'inner world' (or subjective realm). The outer world is populated by things such as mountains, cars, governments, and other things whose existence is independent of subjects. The 'inner world' is populated by such things as sensations, 'visual impressions,' and retinal images, which depend for their existence on the existence of subjects. Frege's problem, concerning perception, is to understand how a sensation, which is part of the inner world, can give rise to perception of the outer world. Sensations arise in us by the effect of the outer world on our senses, but they do not "disclose the outer world to us"; for our sensations are individual to us and vary independently of the facts of the outer world. And yet we clearly do each possess the capacity to perceive the outer world on the basis of sense impressions, which enables us to "move about in the same outer world." Frege's answer, in the foregoing

passage, is that sensory capacities must be supplemented by a proposition of the sort to which one is related in thought. (In fact, Frege holds the stronger claim that sensation must be supplemented by a propositional *thought*; that a judgment is "what sensation lacks to make perception possible.") Frege reaches this conclusion through two further claims. The first is that objectivity requires a notion of error; truth or falsity, correctness or incorrectness. Frege's point is not merely the weak claim that the notion of a truth-condition provides us with a useful theoretical device for capturing the 'aboutness' or intentionality of a thought. Rather, he regards the notion of truth as *constitutive* of objectivity. As Adrian Cussins expresses the connection between these notions,

> A start on the objectivity of content is this: that the content's referent is given as public, as something which is, in principle at least, equally available to any subjective point of view. A sign of this objectivity is that the content can be incorrect: If the referent is given as a public object, then it is always possible that the subject is wrong about the object, even where the referent is the subject oneself. What we are after is a metaphysical distance between subject and object, a distance which makes intelligible the subject's being wrong (and therefore also being right) about the object; which provides for the possibility of truth. (1992: 660)

The point of the normativity is that for thought to be *about* mind-independent objects " the whole environment, in which I am supposed to move and to act," one must be capable of taking them as such; that is, one must be capable of taking objects "for another person as

for myself"; and able "to distinguish that of which I [am] the bearer from that of which I [am] not" (ibid., 303). Frege's second, distinct claim concerns what we are committed to if we accept this connection between objectivity and truth: 'a third realm must be recognized'; a realm of objects which are neither perceptible objects nor subjective states or ideas but *propositions*; sui generis, non-sensible objects, existing independently of minds and outside of space and time, which are the fundamental bearers of truth and falsity. Combining these claims, the senses must be supplemented by precisely those kinds of abstract objects by means of which we *judge*.

We can regard Travis's view of experience as a product of both what he rejects in Frege's picture and what he accepts in it. What he emphatically rejects is Frege's conclusion that perceptual awareness is given by an abstract proposition. Travis thinks that it is the very essence of perception, which distinguishes it from judgement, that it is a *direct* relation to objects. One reason for this, as we have seen, is that he thinks this is required for experience to bear on belief. But there are also reasons of a more general and intuitive nature which one might think count against Frege's conclusion. Another reason is that this is simply inappropriate to characterizing the relationship between perceptual content and sensory experience. The perception is not like the judgement; the judgement is indirect whereas the experience somehow connects us *directly* to objects in the world. As Tim Crane has put it, "(W)hen an object is perceptually experienced, it is experienced as "there," "given," or "present to the mind" in a way in which it is not in belief, thought and many other mental states and events... perception can only confront what is presently given: in this sense, it seems that you can only see or hear or touch what is there" (Crane and

French: 2017). Since Travis rejects Frege's conclusion, he must reject one or other of the premises which lead to it. As we have seen, Travis follows Frege in asserting an "unbreachable" divide between representational content and sensible objects; between generalities, which are instanced, and history, which does the instancing" (2012: 4). He accepts the identification of correctness conditions with propositions. Consequently, he must reject the necessary connection between objectivity and correctness conditions. For Travis, there can be a kind of objectivity that consists in being related to objects that are not representational contents and that do not essentially have function of being correct or incorrect.

Evans's conception of the nonconceptual also involves a rejection of Frege's conclusion: 'nonconceptual content' has objective significance but it is not a proposition. This is connected to Evans's interest in explaining demonstrative reference. Demonstrative reference involves an "unmediated disposition to treat information from that object as germane to the truth and falsity of thoughts involving that Idea"; this disposition in turn "rests upon certain very fundamental perceptual skills which we possess: the ability to keep track of an object in a visual array, or to follow an instrument in a complex and evolving pattern of sound" (1982: 146). The nonconceptual content of perception explains how demonstrative reference is possible by capturing the particularity of perception. But Evans also agrees with Frege in another important respect: he agrees that objectivity must be captured in terms of norms of *truth* or *correctness*. Informational content, including nonconceptual content, is about the world for Evans because it meets this condition: "the content of this piece of information … can be said to embody information from x. It is with x

71

that the predicative material in the informational state has to be compared if we are to evaluate how successfully the system has worked." Relatedly, informational content also admits the possibility of error: an informational state "may be of an object even though its content fails to fit the object at all well—because of malfunction in the system"; it may also be "of nothing: this will be the case if there was no object which served as input to the informational system when the information was produced." Because Evans accepts this component of Frege's framework, his rejection of a propositional content for perception commits him to *rejecting* the other part of Frege's view: the link between normativity and *propositions.*[20] Intuitively, Evans's proposal that nonconceptual spatial experience is captured by an egocentric framework suggests just such a model. 'Locating *objects* in space' is a kind of object-involving content. The egocentric framework is the abstraction; the part that is not the concreteness of sense. Evans's notion of nonconceptual content suggests something that can be both abstract and concrete at the same time.

Yet in spite of its promise, Evans's characterization ultimately fails to deliver. The problem lies in how he 'specifies' or individuates spatial nonconceptual content in terms of 'patterns of inputs and outputs': sensory registrations and actions, or bodily dispositions. To take the example of hearing a sound as coming from such-and-such a position, the auditory input acquires a (nonconceptual) spatial content for an organism by being linked with behavioural output in a certain ("presumably advantageous") way. As Evans elaborates this point,

---

[20] We should be careful not to confuse the notion of 'informational content' with that of a 'nonconceptual content': communication is part of the informational system and "language embodies conceptual information." But "the senses yield nonconceptual information": informational content in this sense can, for Evans, be a kind of nonconceptual content, and is so in the case of perception.

> The spatial information embodied in auditory perception is specifiable only in a vocabulary whose terms derive their meaning partly from being linked with bodily actions. Even given an irreducibility, it would remain the case that possession of such information is directly manifestable in behaviour issuing from no calculation; it is just that there would be indefinitely many ways in which the manifestation can occur. (1982: 157)

The problem centers on the fact that there is, on this conception of spatial perceptual content, no notion of a representational content, independent of how one would act on it, which can be considered the cause of (or explanation for why) one acted that way. It is the system of inputs and outputs that determine the 'content'. But this is not truly a notion of a state which 'has' content, in the sense that thoughts have content. We think of the contents of thoughts as explaining actions, as being causes or reasons for them. If the representational content just is the action, it cannot be like this for perceptual content. More generally, it is hard to see how Evans's view of perception differs from the view of someone like J.J. Gibson, who appealed to perception as an embodied, active engagement with the environment precisely to show that there was no *need* for internal representations in the case of perception. For Gibson, perception was emphatically to be distinguished from sensation: "perception involves meaning; sensation does not." But this 'meaning' or 'objective significance' is accounted for not by adding a *representational content* but by broadening our view of the inputs that are relevant to perception and their connection to

'organismic action' within an environment. The crux of Gibson's 'ecological theory of perception' is that there exist within varying sensory stimulus, considered holistically, certain invariants which suffice to account for objective perception: "When the senses are considered as perceptual systems, all theories of perception become at one stroke unnecessary. It is no longer a question of how the mind operates on the deliverances of sense, or how past experience can organize the data, or even how the brain can process the inputs of the nerves, but simply how information is picked up" (1966: 319). Perception occupies the space that connects these invariants with 'organismic action': "the useful dimensions of sensitivity are those that specify the environment and the observer's relation to the environment... The sense organs are all capable of motor adjustment" (Gibson, 1963). Gibson's intention in rejecting "theories of perception," as A.P. Costall points out, "was not... a denial of any role for theory in psychology, but an insistence that behavior is subject to lawful description in its own right without appeal to "underlying" structures, be they mental, neurological, or quasi-neurological" (1984: 109). In other words, Gibson's position challenged the representational theory of perception. But if Evans's view so closely approximates Gibson's, what reason do we have to continue to call this 'bodily-' or 'actional-' specified meaning a form of *content*?[21]

---

[21] There is also a further issue with Evans's view which makes his notion of nonconceptual content weak, namely his equivocation over the question whether nonconceptual content is really a kind of *personal level content*. Evans insists that nonconceptual perceptual states must be supplemented by concepts in order to become conscious; hence, in order to become part of how *the subject* represents the world as being. As he clarifies this claim: "I do not mean to suggest that only those information-bearing aspects of the sensory input for which the subject has concepts can figure in a report of his experience. It is not necessary, for example, that the subject possess the egocentric concept `to the right' if he is to be able to have the experience of a sound as being to the right. I am not requiring that the content of conscious experience itself be conceptual content. All I am requiring for conscious experience is that the subject exercise some concepts—have some thoughts—and that the content of those thoughts should depend systematically upon the informational properties of the input" (ibid*.,* 159). In his early work on nonconceptual content, Christopher Peacocke

These observations cast new light on McDowell's accusation that Evans is guilty of 'fraudulent labelling':

Evans's position has a deceptively innocent look. It can seem obvious that a possessor of one piece of representational content, whether conceptual or not, can stand in rational relations... to a possessor of another. But ... it is another case of fraudulent labelling to use the word "content" for something we can even so take experience to have, in such a way that reason-constituting relations can intelligibly hold between experiences and judgements. The label serves to mask the fact that the relations between experiences and judgements are being conceived to meet inconsistent demands: to be such as to fit experiences to be reasons for judgements, whilst being outside the reach of rational inquiry. (McDowell, 1994: 53)

McDowell's charge against Evans here has two parts to it. One is the claim that Evans's 'nonconceptual content' is not what it is presented as being, a representational content that is assessable as incorrect or incorrect in the way that a judgement is. Rather, it is something more like a 'given'; a state of awareness that is somehow imbued with objective significance without being normative. The second part is to challenge what Evans's 'nonconceptual content' turns out to be. We have seen that the first objection, at the very least, is justified. But now, if we accept these points, it follows that one of the key objections the conceptual representationalist has against the nonconceptual representationalist is

follows Evans in holding that the nonconceptual content of perception cannot operate autonomously of conceptual capacities (2003); under pressure from Bermudez (1995), amongst others, Peacocke has since changed his position.

purely verbal. McDowell's objection can't possibly be against a nonconceptual content for perception—because the only view of 'nonconceptual content' he considers is Evans's, which by his own admission is not really a nonconceptual *content* at all. Instead, the substantive disagreement concerning perceptual justification is the one the conceptualist has with the naive realist. At least with regards to the issue of perceptual justification, the disagreement is not between two views about perceptual *content*; rather it is between two views about the structure of perceptual objectivity: one that regards perception as an indirect relation to objects (via a conceptual content which functions to 'match' the world according to its conceptual constituents); the other that regards perception as a non-representationally direct relation to objects.[22]

A cursory examination of the literature in defense of nonconceptual content reveals that this problem goes beyond Evans. Consider Tim Crane's 2012 paper, 'The Given,' which argues against a propositional account of perceptual experience.[23] According to Crane, propositions cannot capture the specific, concrete nature of perception: "it is not easy to

---

[22] Byrne's charge of incoherence is based on a commitment he attributes to the nonconceptual representationalist: that one can think *about* the content of one's experience but not *with* the contents of one's experience. Thinking about a perceptual content, *p*, is having a thought whose content makes reference to a perception which is partly individuated by *p* (e.g. a belief that *S has an experience as of a red object*). Thinking with *p* is having a thought which has *p* as its content or as part of its content (e.g. a belief *that that object is red*). According to the nonconceptualist, one can't think with *p* because thoughts have a conceptual content; hence, the nonconceptual contents of experience are not the sort of thing that could be the content of a thought. But then it is hard to see how reflection on experience could possibly lead one reasonably to suspect that its content is nonconceptual: "One starts with a thought like 'It appears to me that my environment is thus-and-so,' and ends with something like 'So I suppose the content of my experience is rich/perspectival/phenomenal/nonconceptual....' If the premise is to have any bearing on the conclusion, the content one ends up thinking about must be the content one started thinking with, in which case no sensible conclusion can be that the content is nonconceptual" (2005: 24).

[23] Propositional models of experience may involve a commitment to conceptualism but they need not do so. That is, whilst some philosophers who equate perceptual contents with propositions also identify these as conceptual (Frege, 1918/1956; Byrne, 2001), there are also some who adopt a neutral view of the propositional contents of experience, leaving open whether the propositionality of perception is also a conceptuality (Searle, 1983; Seigel, 2010). On the other hand, rejecting a propositional model of experience entails the rejection of conceptualism.

make literal sense of the idea that what we take in in experience is what we can judge. When I judge, because of what I can see, that the pig is under the oak, this is something which in a certain way, abstracts from the real presence of the pig there" (ibid., 16). For "the content of the judgement can outlive the experience, it can be the content of others' judgement, things can follow from it (for example, that something is underneath the oak)"; and so likewise if experience has a propositional content, then this too would be something that can be shared between different subjects; and "which abstracts from the particularity of the subject's own condition." Perceptual content, if such there is, must be "something concrete"; something that is 'prior' to the propositional 'descriptions' we might (validly or truly) apply to it in order to characterize it. Thus, according to Crane, experiences have non-propositional content, "in the sense that their fundamental way of representing the world is non-propositional." But consider how Crane describes this form of 'nonconceptual content':

> The distinction between a concrete, particular act with what I am calling its 'real' content, and the abstract, 'general' content which can be assigned to the act, should be fundamental to the theory of intentionality. The propositional content which can be assigned to an act is 'abstract' not just in the sense that it is an object which has no spatiotemporal location, but in the sense that it 'abstracts' from some of the concrete reality of the experiential episode. If the propositional content of the act is something that can be shared between different subjects, or something that can be shared in different acts of the same subject, then it is something which abstracts

from the particularity of the subject's own condition. The real content, however, is

unrepeatable because essentially linked to the state and time of the act's occurrence,

and specific to its bearer. In this sense of content, no-one other than me can have

mental episodes with the content of my mental episodes. Someone could have a very

similar experience, of similar things, or even an experience which seemed exactly

the same. But that would be to generalise across different concrete experiences, and

describe the sense in which they are the same. Describing is relating the experience

to a propositional content. But this description is an attempt to capture some aspect

of how the representation represents the world. It—the description—is not the

representation itself. (2012: 19)

'Nonconceptual content', on Crane's usage, is not a form of representational content in the

standard use of that term. To be non-propositional, for Crane, is to be concrete; hence, not

assessable for accuracy. Experiences can come to be associated with a correctness

condition, through assigning a conceptualized 'interpretation' to them—but "the nature of

that experience itself as a representation is what is described by the correctness condition;

it does not consist in the assignment of a correctness condition." We must not confuse this

"theoretical, external description of the state" for "what it is in itself" (ibid., 22). In rejecting

a propositional content for perception, then, Crane is rejecting the correspondence model

of perceptual awareness. Although he uses the term 'content,' Crane is endorsing the same

part of Frege's argument that Travis endorses: *does not impose correctness conditions*

unless they are supplemented by a proposition which specifies principles of individuation.

Crane anticipates this objection to his use of the term 'content':

Some philosophers will say that the only way in which something can be a representation is if it is interpreted as such, if it can be assigned a correctness condition. I assume here that this is false. There can be an intrinsic form of representation which is not simply a matter of the assignment of a correctness condition by an interpreter, of something being described as representing something. An experience is, on this view, such a representation. What the assignment of a correctness condition does is to specify one or more of the many ways in which the experience represents the world. But the nature of that experience itself as a representation is what is described by the correctness condition; it does not consist in the assignment of a correctness condition. (ibid., 20)

For Crane, what it means to say that experiences—these particular, conscious episodes—have (non-propositional) content is that since "they are a case of something being given or conveyed to the subject" (ibid., 19). Here, Crane is endorsing Travis's point: we should reject the idea of a necessary connection between objectivity and truth. But the term, 'content,' is being used in this context to refer to a fundamentally different kind of thing than the kind of thing that is picked out by 'the contents of thought.' The similarity is only superficial. At the level of ontology, these are two different things.

My point in positioning Travis and Evans in relation to Frege was to bring out two conceptually distinct notions of 'the nonconceptual' which are on the table. One of the

conclusions we can draw from this discussion is that it is hard to assess the arguments against nonconceptual content. For we have seen that many times the view that is being challenged under that name is really a naive realist view.

1.5 Two Perspectives on 'Frege's Line'

Although McDowell's epistemological objection to nonconceptual content turns out not to be directed against nonconceptual content after all, I do think there is a substantive disagreement McDowell has with the nonconceptual representationalist. He objects to it on a priori grounds. For McDowell, both 'nonconceptual content' and 'the given' are myths, but they are different kinds of myth. Travis puts the point in more straightforward terms:

> Some have spoken of something called 'non-conceptual (representational) content.'
> On the present idea of the conceptual, that idea makes no sense. Representing as so
> is essentially conceptual. (2007: 232)

This brings me to the other point I want to make. When we recognize the naive realist as part of the nonconceptualism debate, we discover that there are two ways of opposing nonconceptual content. For both McDowell and Travis, the middle way that is suggested by the term 'nonconceptual content' does not exist: representational content is *essentially* conceptual. They also agree, then, about what our options are concerning perception: either perceptual experience is conceptual, or it is not a form of representation at all. Moreover, in both cases this assumption plays a key role in justifying the conclusions they

want to draw: it is in large part the fact that perception cannot be one way that means it must be the other. What distinguishes the two is only which of these sides they regard as unacceptable.

For McDowell, the dichotomy I have described is expressed through the Sellarsian dichotomy of logical spaces. One is 'the space of nature,' that is, the realm of physical things, governed by causal laws. The other is 'the space of reasons,' that is, the realm of representational content, governed by normative laws. 'Normative,' in this context, means subject to norms of truth or correctness:

> To make sense of the idea of a mental state's or episode's being directed towards the world, in the way in which, say, a belief or judgement is, we need to put the state or episode in a normative context. A belief or judgement to the effect that things are thus and so—a belief or judgement whose content (as we say) is that things are thus and so—must be a posture or stance that is correctly or incorrectly adopted according to whether or not things are indeed thus and so… This relation is normative, then, in this sense: thinking that aims at judgement, or at the fixation of belief, is answerable to the world—to how things are—for whether or not it is correctly executed. (1994: xi)

To characterize perception as a 'given' is to place perception in the 'space of nature.''The given is alternately characterized as: a "bare presence"; something subjective and incapable of being shared ("only one person could be the subject to whom a particular bit of the Given

is given"); 'bits of experiential intake.'[24] It is introduced to fill this role as an external constraint on our freedom to deploy our empirical concepts: the space of reasons, the space of justification or warrants, extends more widely than the conceptual sphere. The extra extent of the space of reasons is supposed to allow it to incorporate nonconceptual impacts from outside the realm of thought. In the attempt to introduce an external constraint by going outside of the space of reasons, we abandon the possibility that perception can bear on what we are to think.

For Travis, the dichotomy is expressed through Frege's idea that the domain of sensible things and the domain of representational contents are separated by a *sharp line*. That something belongs on one side of the line means that it does not belong on the other: that something belongs to the realm of representational content means it does not belong to the realm of sensible particulars; that something belongs to the realm of sensible particulars means it does not belong to the realm of representational content. (It is in this sense that what is perceivable is excluded from "the domain of things for which truth can come into question at all" (Frege 1918/1956: 61)). Travis's argument for this conclusion is in effect to turn McDowell's desiderata against him. That is, let us grant that it is a condition on any theory of experience that it give an account of how experience can 'bearing on what one is to think,' by 'settling the question whether P, or making P likely, or being evidence for P.' Then we must reject the view that experience is conceptual. For a conceptual generality cannot itself settle the question of whether something in the world fits it or correctly falls under it. Rather, it is the world itself that settles that question; it is how

---

[24] See particularly his discussion of Wittgenstein (ibid.,19-23) for this point.

things are in the *world,* and not how things are with us, that determines whether our concepts are properly applied. The only way experience could possibly bear on conceptual generalities is if experience were a relation to the *nonconceptual*. Hence, "there precisely must be rational relations between the conceptual (what satisfies the condition) and something else if we are to make sense of experience bearing on what one is to think" (ibid., 221).

If we take the arguments of McDowell and Travis in conjunction with one another, we can generate a dilemma. For we have three mutually inconsistent claims:

(P1)   Either perceptual experience is a relation to a conceptual content or perceptual experience is a non-representationally direct relation to mind-independent objects and/or properties

(P2)   Experience is not a non-representationally direct relation to mind-independent objects and/or properties

(P3)   Experience is not a relation to a conceptual content.

It cannot be the case that (P1), (P2) and (P3) are all true. At least one of these claims must be false. But which should we reject? Travis, of course, does not regard this as a dilemma. He thinks it is clear that we must reject (P2). His point is that, in order to be a reason for belief, perception must be something concrete, particular; something that can satisfy a perceptual generality. But notice that this is not an answer to McDowell's question, so much as a compelling countervailing consideration. In effect, Travis's move in this game is

to introduce new considerations concerning what would be required for perceptual experience to be a *reason* for belief. And I, for one, find Travis's case for why we should abandon the view that experience has a conceptual content compelling. But notice that it is consistent with finding Travis's case against the conceptualist compelling that I also find *McDowell's case against the given compelling*. But if these are supposed to be our only two options, and we are presented with two sets of equally compelling considerations which preclude our adopting either, then we have a problem.

The obvious way out of this dilemma, of course, would be to reject the claim that the distinction between these two 'realms' is exhaustive. And one way to do that would be to reject the claim that the distinction is *exclusive*: to reject the idea that there really are two self-contained *spaces* here; or that there really is one *line*. "What has visual, auditory, or spatial, properties," Travis tells us, can represent,

> only insofar as an intention attaches to it. That is, only insofar as it is to be taken as representing in a particular way (as opposed to others in which, for all its perceivable features, it might). If that painting represented Chartres cathedral as looking thus and so, that is in part because that blue patch in it, on the image of a wall, is to be taken as mattering in a certain way to what one would see in viewing the cathedral if it were as represented (whether, say, the wall would be coloured, or merely in shadow). We need a way of taking what is perceivable for the way things are (the world) to matter in some determinate way to whether things are as it represents them. An intention's function here would be to make the world matter in

one definite way rather than others. For a question of truth to arise just is for the

world to matter in some such way. (2007: 230)

These observations are offered as a (partial) justification for distinguishing sensible objects

from propositions. A sensible object, taken by itself, is not the sort of thing for which it

makes sense to ask, 'Is it true that…?' It requires supplement. According to Travis, and

likewise according to Frege and McDowell, the addition of an intention must take the form

of supplement by a *proposition*. What is added must be a *fully* abstract object, containing

nothing sensible. But what argument is given for this? Why couldn't the requirement be

met by supplementing with something abstract that is not yet a self-contained

representational content? Earlier, we considered Evans's idea of an ego-centric

'framework.' I believe (to anticipate the next chapter) that it is also an idea we find in

Kant's view of intuition as composed as a 'form' and a sensory matter. Such an object, if it

existed, would be an entity straddling Frege's line: a genuine representational content that

is neither (fully) abstract nor (fully) sensation, but both of these at once. To pursue this

suggestion would be to pursue McDowell's idea of perception as a fundamental unity

between sensation and representational content *at the level of the content itself*: the content

is a hybrid of the physical and the abstract, the physical part is the sensory registration; the

abstract part is the frame of reference, the coordinate space.

The advantage of thinking about representational content in this way, of course, is

that there is a role for *both* the abstract and the concrete. We do not have to make an

awkward choice. This is one way of thinking about what 'nonconceptual content' might be:

not like a sentence or like a picture but something more like a graph. To my mind, the difficulties we have considered in this section make this suggestion worth pursuing.

1.6 Conclusion

In this chapter I have argued that the nonconceptualism debate is more complex than has generally been recognized. Granting that perception is distinct from belief, the question that is at issue in that debate is whether we also need to acknowledge a difference in the *objects* of perception—something other than the conceptual contents which are the direct objects of our beliefs about the world. My main aim has been to show that this question admits of more varied answers than have generally been recognized. Traditionally, the debate is presented as a disagreement between two parties: those who follow Gareth Evans in holding that an adequate account of the differences between perception and judgement requires us to recognize a 'nonconceptual content' for perception; and those who follow John McDowell in maintaining that the deep differences between perception and thought must respect a fundamental continuity of content between the two different kinds of states. I have argued that we must recognize as a third position, opposed to both Evans and McDowell, the views of philosophers such as Charles Travis and Bill Brewer for whom the 'nonconceptual' character of perceptual experience consists in its being a non-representationally direct relation between subjects and the constituents of the mind-independent world. The result of this expanded perspective is that the role of concepts in perception—whether there is, or could be, an form of objectivity which did not use the machinery of concepts—turns out to be only part of what's at issue in the

nonconceptualism debate. For it is also and equally a disagreement about the nature of 'the nonconceptual.' Is this alternative to conceptual objectivity carved out from within the existing framework of mental content, as Evans suggests? Or is it instead the case that, in order to make space for the nonconceptual, we must go outside of framework; throwing out the orthodoxy that objectivity must be representational in nature?

The fact that there are two distinct conceptions of 'the nonconceptual' is important because it matters, for the purposes of assessing the disagreement between the conceptualist and the nonconceptualist, which of these conceptions the conceptualist is addressing themselves *to*. McDowell's epistemic argument, which criticizes a picture of perception as a brute, uninterpreted 'Given,' is primarily an argument against a certain conception of the nonconceptual. When McDowell charges Evans with 'fraudulently labelling' the given as a form of content, he is equating Evans's view with that of the naive realist; his arguments against Evans are predicated on that equivocation. Thus, his complaint against each side is different: whilst the 'direct' conception of the nonconceptual is inadequate, the 'content' conception of the nonconceptual is simply *incoherent*; to subtract from the contents of thought their 'conceptuality' is to be left with something to which the norms of correctness or truth which govern representational content can no longer apply. Travis and Brewer share, and are partly motivated by, this conviction that representational content is *essentially conceptual*. They regard theirs is the only coherent option for one who acknowledges the need for the nonconceptual in perception. As I argued in section 1.5, the exclusion of a 'middle way' is at the heart of a dilemma concerning the role of perception in justifying beliefs about the external world.

What emerges from this analysis is a deeper understanding of what has motivated the existing debate and an indication of where progress lies. Much of the action rests on an a priori rejection of nonconceptual content, an in-principle exclusion of a third way between the position of the conceptualist and that of the naive realist. As I argued in section 1.4, Evans bears much of the responsibility for this. His conception of nonconceptual information as a kind of content demarcates something that, as a form of content, has one foot in the world of thought despite being outside of it by being concrete and embodied. Yet Evans fails to explain how it is possible to subtract conceptual abstraction and be left with something that would qualify it as a kind of personal-level 'representational content.' Evans's lack of care to this question lends legitimacy to McDowell's charge of fraudulence; and many of those who have endorsed the 'content' conception of the nonconceptual following Evans have cemented his mistake (Crane, 2009, 2013; Dretske,). But I do not think we can conclude from the fact that the challenge has not been met that it *cannot* be met. Indeed there are substantial enticements to try to meet it, namely the promise of resolving the dilemma that emerges from the exclusion of the middle way. Moving forward, a lot hinges on whether the friend of nonconceptual content can succeed where Evans failed, by explaining in concrete terms how representational content can perform the same function as conceptual content in the absence of concepts.

In summarizing the moral of this chapter, and its implications for the chapters to follow, I should like to quote Hartry Field on the difficulties with philosophical discussions of mental content. In 'Stalnacker on Intentionality,' Field urges the need for clarity in

distinguishing the 'rough and ready senses in which a mental state can have content' from constitutive claims about contentful mental states:

> (O)ur relatively ordinary assertions of contentfulness and of sameness of content seem highly context-dependent—especially assertions of sameness of content between the mental states of different agents. (Even more especially, when those agents do not share a language; still more, when they don't even belong to the same species.) Any view according to which we are to assign entities to mental states that are to serve as their contents (and are then to define having of-content and sameness-of-content in terms of such assigned entities) is clearly ladening our ordinary talk of having-of-content and sameness-of-content with a substantial body of theory; and my ... point is that the nature of the theory and the motivation for introducing it deserve serious discussion. (2001: 83)

Our discussion in this chapter illustrates in a specific way the confusion that results from a failure to give attention to precisely those issues Field recommends: the nature of the theory of (nonconceptual) *content*, as opposed to something that is 'content' only in a rough and ready sense. To advance that debate, the proponent of nonconceptual content must take the 'content' side of the equation as seriously as the 'nonconceptual' side. How should we think of the 'nonconceptual' so that what is subtracted from conceptual content is still meaningfully a form of *representational content*? What kinds of connections exist between 'content' and 'concept'; and how far does the legitimacy of a priori theorizing extend in this

domain? What kinds of considerations would provide evidence for the 'the nonconceptual' in experience that was not simultaneously evidence that experience is not a form of representation at all? What, beyond the 'aboutness' of perception, is perceptual content required to explain? The chapters that follow aim to answer these questions. The next chapter takes up the theoretical question of what 'nonconceptual content' might be by carving out a 'middle way' between conceptualist abstraction and the pure particularity of the naive realist view. Chapter three takes up the question of what would count as a criterion for nonconceptual content through a consideration of Evans's Generality Constraint. Finally, chapter four proposes a positive argument for the view of nonconceptual content defended earlier which anchors this in a broad picture of how representational contents contribute to explanations of perceptual capacities.

2. CLASSIFICATION WITHOUT CONCEPTS

Nothing is more difficult than confronting concepts without prejudice—For a

prejudice is a system, and hence a form of understanding, though not the right one.

(Wittgenstein, M.S. 136: 18. Quoted in Schulte, 2014: 21)

2.1 Introduction

Compare these two characterizations of concepts:

Concepts seem to be the very stuff of which cognitions are made. At any rate,

cognitive states like beliefs and preferences, with which many of us hope to explain

behavior, seem to involve relations between agents and, roughly speaking,

conceptual contents. … Concepts … [provide] the commonalities between different

contents, the links between different cognitive states that are 'about the same thing'

(cf., the atoms and their structures as the basis for the Periodic Table)... Thus, a

theory of concepts fulfilling this function ought to provide a basis for characterizing

the relations, e.g., between a belief in some generalization ([Cars can skid in

puddles]) and beliefs in its instances ([My car can skid in that puddle]) (Rey, 1982:

242)

Once we consider the role our beliefs play in reasoning, then it starts to become

clear why their contents need constituents... A thinker who believes that a is F, and

that b is F, and that a is not b will be disposed to believe that at least two things are F. [To explain the validity of this inference] (I)t is essential that the same (type of) 'part' of the content should occur in the two states, and that 'it should have exactly the same meaning in both cases'—i.e. the parts should be tokens of the same semantic type. ... The states must both contain F as a part. I say that it is only a terminological variant of this to say that they must contain the concept F. (Crane, 1992: 12)

The main idea expressed in each of these passages is that concepts essentially have the function of classifying particulars as being of some type. They are representations capable of "holding for" or being true of indefinitely many singular terms, in which case those entities are thought to be *of the same type* or *belonging to the same class*. The concept of a tree represents a way for a particular to be which can be true of indefinitely many singular terms. It is because they have this function that ascribing a conceptual structure to thought allows us to explain the "unity" or "connectivity" that exists amongst a subject's body of thoughts. The distinct thoughts that *this object is a tree* and a judgement that *all trees are plants* are partly about the same thing because they both contain the same concept*: is a* tree. Notice, however, that the second passage contains a further claim that is not part of the first. Crane asserts an equivalency between the statements "a thought-content contains F as a part" and "a thought-content contains the concept F." He thereby suggests that as well as possessing the classification function essentially, concepts are the *unique* bearers of this property. To perform the function a concept performs in a thought content just is what

it is for something to be a concept. Crane's position, which treats concepts as synonymous with classification, reflects a widespread view in philosophy of mind. As Eleanor Rosch has noted, "since at least the nineteenth century, it has been common to refer to the cognitive or mental aspect of categories as concepts" (Rosch, 1999: 61). My aim in this chapter is to argue against it. All concepts classify things as being some way. But not everything which classifies things as being some way need be a concept. Alongside conceptual classification, we must also acknowledge the possibility of nonconceptual classification.

My investigation starts from the recognition that, if classification and conceptualization are indeed coextensive notions, this is a substantive claim and not something that is true simply by definition. When we think about what is required for classification, the key notion seems to be that of a structured representational content: a representational content composed of 'term-sized' elements that can common to distinct, correctness-conditional wholes. When we think about what is involved in something being a concept—a constituent of thoughts, capable of supporting rational and linguistic competencies—we tend to be thinking of constituents with a 'word-like' or *atomic* structure. The concept of tree is atomic inasmuch as it refers to trees *and nothing else*. But an entity's being a constituent does not trivially entail its being an *atomic* constituent. So an entity's having the classification-function does not trivially entail its being a concept. The move from the claim that the conceptual constituents of thoughts are essentially ways of classifying particulars to the claim that anything that performs the function of classifying objects *is* a concept must be justified by arguments which demonstrate a constitutive connection between a content's having constituents and its having *atomic* constituents.

There are, I believe, broadly two sorts of considerations which explain the monist's stance. One reason why philosophers are monists is they believe there is a conceptual relation between attribution and predication: all classification is conceptual because any representational content that has an attributive structure must have a predicative structure. I identify two arguments for this claim. The first, which I call the 'argument from mereology,' is exemplified in Fred Dretske's comment:

> (u)ntil information has been lost, or discarded, an information processing system has failed to treat different things as essentially the same. It has failed to classify or categorize, failed to generalize, failed to "recognize" the input as being an instance (token) of a more general type. (1981: 29)

An attributive, for instance x *is a leaf*, is a content capable of contributing the same semantic value to different complex representational contents, thereby grouping different things as being *of the same type*. But this is possible do this only if it 'leaves out' determining information (for example, only by representing the property of being a tree *and nothing more specific*. The second argument, 'the argument from the nature of propositions,' is one we have already encountered, in the previous chapter, in Charles Travis's discussion of 'Frege's Line.' Travis's identification of content with conceptual content is based in the idea that, since a representational content is a proposition—a context-independent, abstract object without a location in space or time —the constituents

of which propositions are composed must be abstract and context-independent in the same way.

I shall argue that all of these justifications fall short. Ultimately, these criticisms will provide the basis for a positive argument in defense of the possibility of nonconceptual classification. Whereas monism is propelled by abstract theorizing about what classification *must* be, the roots of the pluralist view I propose lie in an open-minded investigation of the diversity of structures that actually exist in mind and world. At its basis lies the recognition of a distinctive kind of holism present in magnitudes—including colors, spatial properties, sounds, and in the device of a scale by which we represent such properties. In the words of Frege, "a magnitude is not something all by itself, but only insofar as it belongs in a system with other magnitudes of the same kind." As Wittgenstein describes, the representational device of a scale possesses the same kind of holism inasmuch as,

> when I lay a yardstick against a spatial object, I apply all the graduation marks simultaneously. It's not the individual graduation marks that are applied, it's the whole scale. If I know that the object reaches up to the tenth graduation mark, I also know immediately that it doesn't reach the eleventh, twelfth, etc. The assertions telling me the length of an object form a system, a system of propositions. It's such a whole system which is compared with reality, not a single proposition. If, for instance, I say such and such a point in the visual field is blue, I not only know that, I

also know that the point isn't green, isn't red, isn't yellow etc. I have simultaneously

applied the whole colour scale. (1998: 317)

Using the example of visual perception, I show how this kind of structure forms the basis

for a *non-propositional*, *non-atomic* form of attribution. In 'spatial attribution,' a sensory

magnitude is represented as being of a certain type (for instance, a given spatial

magnitude) by locating it in the part of the 'frame of reference' which refers to that type.

The form of attribution non-propositional because it contains a fundamentally subjective

or context-dependent aspect. For example, a visual spatial perception which represents an

edge and its length from two different angles will appear different, despite representing the

edge as having the same property. In one sense, the edge looks the same (namely with

regards to its length). But in another sense it looks different. The presentation of an

objective type is always accompanied by a subjective aspect. The form of attribution is non-

atomic because, in keeping with Wittgenstein's observations, the representation of any

given piece of information is always accompanied by information concerning its relation to

the system of properties. For example, one cannot visually represent an edge without

representing its length, shape, and orientation. In addition, the freeing of attribution, in

sentence-like structures, from a role in context-bound singling-out of particulars marks

what Tyler Burge calls "a subtle kind of freedom from the here and now." "Pure attribution,

including conceptual attributives, marks a capacity to separate attribution, a constitutive

element in any representational perspective, from its role in guiding contextual singling-

out of particulars that have a causal impact on the individual and the individual's

perspective" (2010: 542). The claim that any form of attribution must be atomic is false. Moreover, the contrast between atomic and non-atomic forms of classification give rise to theoretically significant differences in what one can do with those representations. The freedom inherent in atomic representations is a crucial precondition for the development of capacities for propositional inference, and for the representation of abstract, non-perceptual qualities.

A notable implication of this conclusion is to provide us with a theory of nonconceptual content capable of meeting the requirements discussed in the previous chapter. It is interesting that, whilst the question whether there could be classification without concepts bears obvious connections to the question whether there could be *content* without concepts, in practice the oppositions of monist/pluralist and conceptualist/nonconceptualist have not generally lined up neatly with one another. Certainly, monism is a motivating assumption for those who hold that contents are "essentially conceptual." But the idea that structures of classification are coextensive with structures of conceptualization is also taken for granted by many—if not most—philosophers sympathetic to a nonconceptual content. In this respect, I believe, they are doubly at fault: not only for their assumption that classification without concepts is impossible, but also for having willingly surrendered what I regard as the key battleground in the fight over nonconceptual content. The best chance for a compelling and theoretically viable account of non-conceptual content lies in the possibility of classification without concepts. Looking beyond the nonconceptualism debate, the conclusion of this chapter also

promises to shed light on other important questions, such as the essential unity of representational content.

2.2 Monism

That concepts are ways of classifying objects is about as close to a universally accepted idea as one is likely to find in theorizing about concepts. What marks the position of the monist about classification is that they also endorse a necessary connection in the opposite direction: classification occurs only through concepts. A monistic view of classification is consistent with a range of different views about the nature of representation more generally. Some philosophers who are monists about classification are also monists about representational content: there is no classification without concepts; and there is no representation without concepts either (Byrne, 2005; Fodor, 2007, 2011; Travis, 2007; Brewer, 2006). For some of these philosophers, moreover, the two forms of monism are connected. But it is important to note that many friends of nonconceptual content have also been monists about classification. Philosophers including Richard Heck (2000, 2007), Fred Dretske (1981), Tim Crane (1992, 2008) and Roblin Meeks (2006) claim that although classification is unique to conceptual content, there are forms of representational content which are nonconceptual because they do not involve classification.[25]

My aim in this section is to clarify what is involved in this claim. To have a meaningful discussion, we need to be clear about what we mean when we talk about

---

[25] Richard Heck endorses both claims when he assumes as a constraint on an account of nonconceptual content that this be unstructured content: "The thesis that perceptual content is nonconceptual, as I understand it, implies… [that] The content of perceptual states would … have to lack the conceptual articulation characteristic of Thoughts (as, say, a set of possible worlds does)."

'classification' and 'concepts'—especially since these terms are used with a variety of different meanings by philosophers. What I want to bring out in particular is the sense in which standard philosophical commitments concerning the natures of concepts. The result of this is that monism amounts to a substantive claim about the nature of classification—one which stands in need of defense.

Whilst there are different views about the nature of concepts, most philosophers agree that concepts are somehow associated with elements in propositional structures. Concepts are individuated, either in part or in whole, with elements in structured propositional contents. A proposition is a mind-independent abstract object which determines a truth condition, a condition for being true of a subject matter. A structured proposition is a complex proposition having a part-whole structure that is isomorphic to the syntactic structure of the sentence which best expresses it: we can "distinguish parts in the [proposition] corresponding to parts of a sentence, so that the structure of the sentence serves as an image of the structure of the [proposition]" (Frege, 1923/1963: 1). To illustrate this suggestion in concrete terms, consider the sentence in (S) and its syntactic analysis as given in (S'):

(S)   Croesus is king of Lydia

(S')   [[Croesus]$_{NP}$ [is king of Lydia]$_{VP}$]$_S$

According to the analysis given in (S'), the sentence, "Croesus is king of Lydia" is syntactically complex, being composed of the noun phrase 'Croesus' and the verb phrase 'was king of Lydia,' which are combined through the syntactic operation of predication to form a sentence that is true or false according as the verb phrase is true or false of the

99

object, if any, to which the noun phrase refers. Analogously, the structured proposition that

Croesus is king of Lydia is a complex entity, being composed of a singular term referring to

the individual *Croesus* and a predicate term *is king of Lydia* which are combined through

the syntactic operation of predication to form a proposition that is true or false according

as the predicate term is true or false of the singular term ('Predicate term' may be

construed broadly here to include n-place relations). The two major theories of structured

propositions, Fregean and Russellian, differ in their principles of individuation for

propositional constituents and their views about how the syntactic operation of

predication is realized in propositions.[26] Whilst this dissertation generally follows the

Fregean tradition in assuming that conceptual contents are fine-grained propositions

composed of modes of presentation, these details won't be important for present purposes.

Our concern is with characterizing the structure of conceptual classification, and in this

respect the differences between Fregean and Russellian theories of propositions are much

less important than their shared commitment to a 'sentence-like' structure for

representational contents.

I want to stress two consequences of adopting this framework. First, if we accept

that concepts are associated with elements in structured propositions it follows that the

property of being a concept includes the function of classifying objects. Classification is a

semantic function, by which an object is represented as being of some general *type*, which

other objects could also be. I shall characterize this function through the notion of an

---

[26] On Frege's model, predication is conceived in terms of a functional operation which takes a 'saturated sense' (a mode of presentation of an object) as the input to an 'unsaturated sense' (a mode of presentation of a property) that delivers Truth or Falsity as its output. On the Russellian approach propositions are ordered n-tuples of objects and properties. Propositional unity presents something of a problem for Russell, and over the years Russellians have proposed various ways of resolving this (King, 2007; Collins, 2011).

attributive; that is, a representational content capable by its form of being true of more than one singular terms (Burge, 2005). It is a necessary and sufficient condition for a representational state to classify objects is that it has an 'attributively structured' representational content: a complex content (a is F) composed of a singular element (a) and a general or 'attributive' element (x *is F*). The distinction between singular and general elements marks differences in their functions: a singular element functions to satisfy a general component; whilst a general element functions to be veridical of a singular element. In combination with a singular term, an attributive yields a representational content which is evaluable as correct or incorrect. Different representational contents in which that attributive figures are partly about the same thing. Framed in more intuitive terms, the foregoing proposals are simply a way of characterizing, at a psychological level, a kind of structure we see in the mind-independent world. If we think of a representational state as a kind of 'mental analog' for the worldly state of affairs it represents, then an attributive is the analog for the properties which partly constitute those states of affairs. Consider the belief that Croesus was the ruler of Lydia. On a natural view, the belief is true if and only if there existed an individual, Croesus, who instantiated the property of being ruler of Lydia. These two kinds of things—individuals (or, more generally, particulars) on the one hand and properties on the other—play a mutually supporting role in constituting a state of affairs: a thing is the bearer of a property, and a property is instantiated or possessed by things. It is part of the classical picture that properties are *universals*, being the sorts of things that can be instantiated by numerically distinct things. On this view it is possible for two different individuals, Croesus and Gyges, to exemplify exactly the same

property, being ruler of Lydia. Given these observations, we can think of classification as a kind of representation in which a state represents a state of affairs in a way that ascribes this structure possessed by the state of affairs itself. Just as a property is an entity which is instantiated by many different objects, so the attributive, *is king of Lydia*, is something capable, by virtue of its form, of being *true of* more than one singular term. This allows us to represent commonalities amongst distinct objects; the properties we ascribe to an object in thought is something that other objects could also be thought of as having; in which case we have thoughts which are partly about the same thing. In this way, we are able to capture the connections amongst our thoughts that are central to the explanation of a thinker's inferential and linguistic capacities. A predicate is a kind of attributive: it is a representational content which is capable of being true of more than one object. The fact that something is a concept trivially entails that it has the function of classifying objects.

The second consequence I want to stress is this: if we accept the association of concepts with elements of propositional structures, then 'being a concept' amounts to something more than 'being an element which functions to classify objects.' Bertrand Russell once commented that "every account of structure is relative to certain units which are, for the time being, treated as if they were devoid of structure, but it must never be assumed that these units will not, in another context, have structure which it is important to recognise" (Russell, 1948: 252). Russell's advice is pertinent to our present discussion. We've seen that commitment to structured propositions involves being committed to a whole representational content containing 'units': elements or constituents which contribute their individual contents to the content of the whole. But it also includes a

commitment concerning the structure of those units: both the singular and attributive elements of structured propositions are *word-like* or *atomic* in their form. The significance of this distinction pertains to the different functions these structural properties designate. The structure possessed by a whole representational content determines how the whole performs its function of being correct or incorrect of a subject matter. The atomic structure of concepts, on the other hand, qualifies the way in which concepts refer. Just as the words or phrases that make up a sentence have a meaning of their own, independent of the meaning of the whole sentence which they compose, so an atomic constituent of a proposition has a content of its own, independent of the other parts of the proposition (Crane, 1992: 147). To put it another way, concepts are context-independent or completely abstracted representations of objects and properties: representations of objects and properties *as such*. A concept of F refers to F and nothing else. "the referent is completely determined by the nature of complete senses or concepts (as opposed to indexicals) that the person employs." As Burge puts it, "A trademark of a sense or Fregean thought component is that it… its relation to its referent(s) is atemporal and depends purely on its own nature and the inventory of the world" (Burge, 1979: 238). In assigning an atomic structure to conceptual constituents, then, we are assigning a property over and above what is involved in their having an attributive structure. It is essential to something's being a concept not only that it functions to classify objects, but that it does so by means of a context-independent representation of a property.

I stress this distinction between attribution as such and predication (attribution plus atomism) because it is often overlooked. It is tempting to feel that 'being an

attributive' and 'being a predicate' are two ways of talking about the same thing. Consider Christopher Gauker's claim:

> Iconic representations are never concepts. Iconic representations of a red ball and a blue ball do not represent what the two balls have in common, as the concept ball does. A concept, such as dog, has an argument place. By substituting a representation of a particular object into that place, we can form a whole thought, such as that Fido is a dog. An iconic representation does not have an argument place. There could be a kind of thinking in pictures in which a mental image of a collie did the work of the concept dog in forming the sorts of thoughts we express in English with the word "dog." But in that case, the mental image of a collie would cease to be an iconic representation of a particular collie. (2011: 127)

Gauker identifies concepts with representations having an 'argument place.' It is natural to interpret Gauker's notion of an 'argument place' as equivalent to what I've been calling an attributive: it is a representational constituent which has the function of being true of more than one singular term. By identifying concepts with representations which have an argument place, then, Gauker looks to be adopting a non-standard definition of concepts which makes no reference to their internal structure. The idea that classification is conceptual will be trivially true on this view, simply by virtue of how we are defining the notion of a concept. But I suspect that Gauker is not intentionally omitting the atomic nature of concepts from his account. For notice that, inasmuch as we also tend to associate

the notion of an 'argument place' with a predicative expression or term, this characterization obscures the functional difference that exists being between being an attributive and being a predicate. This characterization smuggles in *atomism* under the radar. This is not to deny that there are some philosophers who may deliberately omit atomism from their definition of a concept, reflecting a principled position about what is important to our characterization of things as 'concepts' (we shall discuss this sort of view in section 2.7). The present point is just that there are many cases in which failure to specify atomism merely reflects a misleading assumption that attribution presupposes atomism, or is somehow inseparable from it.

Based on this discussion, we are now in a position to articulate what is at stake in the truth of monism. According to the monist, conceptualization and classification are coextensive. For any x, if x has the function of classifying objects as being of some type then x is a concept. What emerges from the foregoing observations is that if this is true, it is not trivially true. What is trivially true is that anything which is a concept has the function of classifying objects as part of its nature. For concepts are identified with predicative terms, and predicates are a kind of attributive: they function to be true of indefinitely many distinct singular terms.   But predicates are not *just* attributives, they are *atomic attributives*: attributives which refer to properties in a context-independent way. So the monist is committed to the claim that for any x, if x is an attributive then x is an atomic attributive. But this is not a trivial entailment. Rather, it involves a substantive commitment concerning the nature of classification: attribution is essentially atomic in form; predication

is not merely *a* form through which attribution may be realized but it is the *only* possible form attribution may take.

2.3 The Argument from Mereology

The first argument for monism which we shall consider is based on the idea that the structure possessed by concepts is the same structure any constituent must have in order to be capable of performing the function of an attributive. A natural starting point for this argument is the observation that the word-like structure of concepts is relevant to the way in which they perform their function of classifying or grouping individuals. This idea was implicit in Frege's claim that propositions are "composed of *simple parts*… [which] correspond to the *simple parts of sentences.*" It is made explicit by Locke, who in the third book of the *Essay concerning Human Understanding*, draws a direct connection between the structure of predicate terms and the concepts which they denote:

> Words become general by being made the signs of general ideas: and ideas become general, by separating from them the circumstances of time and place, and any other ideas that may determine them to this or that particular existence. By this way of abstraction they are made capable of representing more individuals than one; each of which having in it a conformity to that abstract idea, is (as we call it) of that sort. (Locke, 1690/1979: III, III, 6)

Locke is describing the formation of predicate terms from concepts ('general ideas'), and the formation of concepts in their turn. A concept (such as X) is formed by a process of abstraction which 'separates from it' any contextual information specific to a time, place or anything else; i.e., by becoming a context-independent representation of a mind-independent property. And through having undergone this process it is "made capable of representing more individuals than one." Their acquiring this atomic structure is what makes them capable of classifying individuals as belonging to a type. The idea that conceptual classification works by analogy with the function of *words* is also emphasized, albeit in a more critical spirit, by Nietzsche. In 'On Truth and Lies in a Nonmoral Sense,' he explains that,

> (a) word becomes a concept insofar as it simultaneously has to fit countless more or less similar cases—which means, purely and simply, cases which are never equal and thus altogether unequal. Every concept arises from the equation of unequal things. Just as it is certain that one leaf is never totally the same as another, so it is certain that the concept "leaf" is formed by arbitrarily discarding these individual differences and by forgetting the distinguishing aspects... We obtain the concept, as we do the form, by overlooking what is individual and actual; whereas nature is acquainted with no forms and no concepts, and likewise with no species, but only with an X which remains inaccessible and undefinable for us. (1873/2006: 117)

Nietzsche starts from a different point than does Locke: he is explaining not how words are formed from concepts but how a word 'becomes' a concept. But the end result of his exposition is essentially the same. The natures of concepts show us how classification works in thought: atomic or 'context-independent' representations of types which abstract by 'discarding' extraneous information. A concept of 'leaf' has the function of classifying different particulars as being of the same type: of equating particular leaves which differ in their particular characteristics. And we obtain the concepts capable of performing this function of 'equating unequal things' by a process of abstraction which 'discards the individual differences' of particular leaves and 'overlooks what is individual and actual.' These observations indicate that the atomic structure of concepts is not incidental to their attributive function. Concepts are capable of being true of multiple, qualitatively distinct objects *because* they leave out what is particular. These observations about how concepts perform their function of classification could be martialled into an argument for monism if it could be shown that what is true of concepts must be true for classification in general.

Fred Dretske's 1981 paper 'Sensation and Perception' suggests a way of making this move. Dretske's commitment to monism is demonstrated in the following passage:

> To describe a process in which a piece of information is converted from analog to digital form is to describe a process that necessarily involves the loss of information. Information is lost because we pass from a structure (the speedometer) of greater informational content to one of lesser information content. Digital conversion is a process in which irrelevant pieces of information are pruned away and discarded.

> Until information has been lost, or discarded, an information processing system has failed to treat different things as essentially the same... [This is] the information-theoretic processes underlying all forms of stimulus generalization, classification, and recognition. (1981: 29-30)

The distinction between digital and analog forms of information carrying presupposed by this passage characterizes two forms in which a signal may carry a certain piece of information. Specifically, a state carries information in digital form if and only if it carries no other information about a not already nested in its being F (as with the sentence "a is F"); whereas a signal carries the information that a is F in analog form if it always includes more specific information (as a photograph of an a that is F might do, for example). For Dretske, this lines up with the distinction between representations that involve classification and those that represent in a way that is *prior* to classification; indeed, the analog-to-digital conversion process contributes to explaining the nature of classification. Notice that Dretske makes two distinguishable claims here: first, that information carried in analog form does not involve stimulus generalization or classification whilst the digital representation does; second, that classification requires the conversion of information into digital form. It is specifically the second of these claims that marks the position of the monist. This process of discarding all information extraneous to F is not only *a* way in which a system may come to classify an object as F; it is the *only* way in which this happens. Dretske does not offer any explicit justification for this further claim. However, a consideration of the context in which it appears helps to clarify his motivations. He

illustrates this point using the example of a system designed to communicate the speed of a

vehicle within a particular range (for example, that it is going at a speed between 15 and 24

or between 25 and 49). The system is illustrated below:
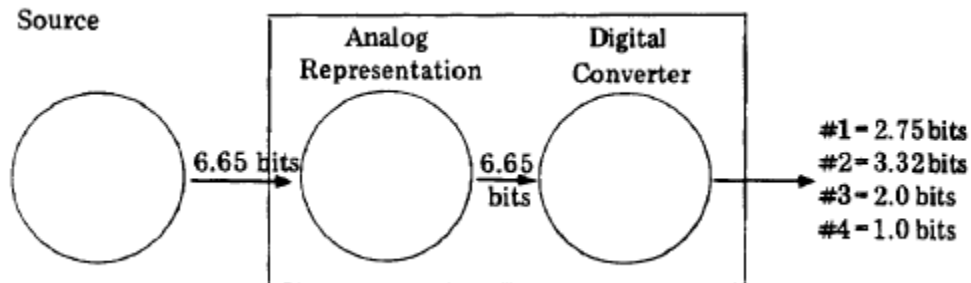
Figure 2. A diagram illustrating a 'digital-conversion system.' The system takes in the registrations of the speedometer in analog form and generates a digital output that classifies the speed within a general range (tones #1-4).

The 'analog component' of this system is the speedometer, which carries all the

information generated by a variable source. Since the source has 100 different possible

states (all equally likely), the speedometer carries 6.65 bits of information about the

source. It carries the information that the vehicle is going, say, 43 mph. The speedometer is

an analog representation of the source because it carries more specific, more determinate

information about the source than is required to control the system's output: not only that

the vehicle is going 43 mph; but also (nested within this piece of information) that the

information that the vehicle is going between 25 and 50 mph. The other part of the system

is a digital converter, which generates, from the information the speedometer carries in

analog form, a representation of the source in digital form: the registrations of the

speedometer are fed into a converter, and (assuming a speed of 43 mph) the third tone is

activated. Since the third tone is activated when, and only when, the vehicle has a speed in the range 25 to 49, this tone carries 2 bits of information about the speed of the vehicle (a reduction of 100 equally likely possibilities to 25). The output of this system is always less, quantitatively, than the input. Although 6.65 bits of information get in, something less than this comes out. The output of this system ignores the difference between 43 mph and 32mph. Both these values are treated as essentially the same. Both activate tone 3. Commenting on the significance of this example, Dretske writes: "This is a form, albeit a fairly primitive form, of stimulus generalization.... The digital converter ... 'throws away' the more specific piece of information [that the vehicle is going 43 mph] and passes along a piece of information (that the vehicle is going somewhere between 25 and 50 mph) that the speedometer carries in analog form are systematically ignored in order to achieve a uniform response to *relevant similarities*." His basis for making this claim is that this is required to achieve a *uniform response to relevant similarities*.[27] That is, in order for the system to represent what is the same across different states, it must *leave out* specific information. The information must be digitized.

We can reconstruct the line of thought behind Dretske's conclusion, I believe, by viewing it in the context of a consideration of parts and wholes. Any representation possessing the function of classifying an object as being of some type must contain a part

---

[27] Tim Crane, citing Frege, claims that in order to explain the validity of inference, "it is essential that the same expression should occur in the two propositions, and that it should have exactly the same meaning in both cases. It must therefore have a meaning of its own, independent of the other parts of the proposition" (1992: 147) The transition from the initial claim to the conclusion is suggestive of the mereological argument. The initial claim is a requirement on classification: the two thoughts must classify the objects in a way that reveals their sameness of type. The second claim draws from this a conclusion that the structures of the constituents must be *atomic*; they must have 'a meaning of their own, independent of the other parts of the proposition.'

which refers to that type and which is capable of being shared in common by different representational wholes. Thus, for example, a state which represents *that a is F* can count as classifying the object, a, as being of the type *F* only if it contains a part which refers to F, and which has an 'argument place' which can be 'filled' by indefinitely many different subject terms {a, b, c, d... }, to generate indefinitely many new representations {a is F, b is F}. For otherwise, there would be nothing in virtue of which different objects could be represented as being of the same type. How does this work in the case of concepts? Suppose I encounter a particular oak leaf lying on the pavement outside my house, and form the belief that that object is a leaf. According to the standard view, I have this thought by virtue of combining a demonstrative representation of the object with my concept of leaf; that is, by bringing the demonstrative under a context-independent representation of the property, being a leaf. This function is made possible through the atomism of the concept of leaf: it refers only to the property of being a leaf and nothing more; it discards all the information that is specific to the character of the individual leaf; and by virtue of being maximally general in this way, my concept is capable of being applied to indefinitely many qualitatively distinct, particular leaves. But now, it is hard to see how there is any other way this result could be achieved. Your concept of a car has to be just as general as the kind, *leaf*. If the concept of a leaf is too specific then all the leaves wouldn't fit under it. It can't include these more specific features that are specific only to some leaves because otherwise it couldn't be true of all leaves; in which case, it wouldn't count as representing the property of being a leaf. The *only* way a component could be legitimately applied to indefinitely many distinct individuals if it didn't 'contain' anything specific to those

individuals; if it 'leaves out' what is different. Anything which has the function of classifying must be atomic.

With the foregoing, we now have an argument for monism. I shall call this the argument from mereology. The claim is that attribution entails atomism because something that did not have an atomic structure could not perform the function of grouping or typing distinct individuals as being *the same*. The result constitutes a substantive conclusion about the nature of classification. All attribution takes the form it takes in conceptual representation. All attribution is predicative in form.

2.4 Attribution without Atomism

In his 1929 essay, 'Some Remarks on Logical Form,' Wittgenstein writes,

> we can only arrive at a correct analysis by what might be called, the logical investigation of the phenomena themselves, i.e., in a certain sense a posteriori, and not by conjecturing about a priori possibilities. One is often tempted to ask from an a priori standpoint: What, after all, can be the only forms of atomic propositions, and to answer, e.g., subject-predicate and relational propositions with two or more terms further, perhaps, propositions relating predicates and relations to one another, and so on. But this, I believe, is mere playing with words. An atomic form cannot be foreseen. And it would be surprising if the actual phenomena had nothing more to teach us about their structure. To such conjectures about the structure of

atomic propositions, we are led by our ordinary language, which uses the subject-predicate and relational form. But in this our language is misleading… (1929: 164)

Wittgenstein's remarks in this paper, together with his (1935) monograph *Philosophical Remarks*, sum up the spirit of his middle-period writings as a reaction against the absolutist tendencies of his own earlier work in the Tractatus concerning the notion of logical form. But they can also be applied to a refutation of the argument from mereology. Certainly, "conjecturing about a priori possibilities" would be an apt description for the methodology of the monist, who proposes conclusions about what *must be the case* for classification to be possible in general based on how classification works in conceptual thought. Wittgenstein's claim is that we should allow our theorizing to be guided by 'a posteriori' investigation of the phenomena themselves; that is, by *looking* at how statements of this kind actually behave. By following his lead, we shall see that what seems valid from the comfort of the philosopher's armchair can be refuted through a consideration of concrete examples.

The specific prompt for Wittgenstein's reexamination of his former view are statements attributing *magnitudes* to objects. These include statements about spatial magnitudes ("this rod has a length of 4 metres"), statements about lightness ("this patch is gray"), statements about colors ("this patch is red") and statements about pitch ("the first tone has a low pitch"). Wittgenstein rejects a subject-predicate form for these kinds of statements on the grounds that their meaning is not adequately captured by a structured proposition composed of a singular term, referring to the objects, and a predicative term referring to the magnitude: [[x] [is 4 meters long]], or [[x] [is gray]], or [[x] [is red]]. One set

of phenomena Wittgenstein identifies as posing a problem for the predicative model concern the 'expectation' that characteristically attaches to statements about magnitudes. For example:

Can absolute silence be confused with inner deafness, meaning having no acquaintance with the concept of sound? If that were so, you couldn't distinguish lacking the sense of hearing from lacking any other sense.

(T)he yardstick must already be applied, I cannot apply it how I like; I can only pick out a point on it. This amounts to saying: if I am surrounded by absolute silence, I cannot join (construct) or not join auditory space on to this silence as I like, i.e., either it is for me 'silence' as opposed to a sound, or the word 'silence' has no meaning for me, i.e., I cannot choose between inner hearing and inner deafness. (1975: 77)

A black colour can become lighter but not louder. But how do these different directions find expression in grammar? Isn't it the same case as my seeing a grey and saying 'I expect this grey to go darker?' How does grammar deal with the distinction between 'lighter' and 'darker'? Or, how can the ruler going from white to black be applied to grey in a particular direction?

If I can only see something black and say it isn't red, how do I know that I am not talking nonsense, i.e., that it could be red, that there is red? Unless red is just another graduation mark on the same scale as black. What is the difference between

'That is not red' and 'That is not abracadabra?' Obviously I need to know that 'black,'

which describes the actual state of affairs (or is used in describing it), is that in

whose place 'red' stands in the description. (ibid., 79)

The 'expectation' here consists in the fact that the meanings of these statements is not

exhausted by what is presented as actual, as being the case; there is also an aspect of

meaning that *is*, so to speak, 'allusive,' involving allusion to facts that 'reach beyond the

present case.' The puzzle this phenomenon presents for subject-predicate form is that *by*

virtue of their semantic atomism, predicate-terms include nothing beyond what is presented

as being the case: the word "grey" makes it seem "as if there is only one point, and how can

I see two directions in that?" The difficulty becomes even more salient in the 'exclusion

problem': how to provide an analysis of the meanings of statements expressing the degree

of a quality so as to capture the fact that certain atomic propositions exclude one another.

Consider a statement about a particular object, a: 'a, is uniformly green and a is uniformly

red." Clearly, there is something wrong with such a claim. The requirement: "It must be

possible for the contradiction to show itself entirely in the symbolism, for if I say of a patch

that it is red and green, it is certainly at most only one of these two, and the contradiction

must be contained in the sense of the two propositions. That two colours won't fit at the

same time in the same place must be contained in their form and the form of space."

[p.107]

Instead, he proposes a fundamentally different analysis of the 'grammar' of

statements about magnitudes in the form of his 'Yardstick Theory of Propositions' (Medina,

2002): which asserts a truth-evaluable proposition by bringing together an entity and a scale. For example, the meaning of the sentence, "That rod is 4 meters long," is captured by the rod's bearing a particular relation to the scale such that its extension within the scale denotes the length that is measured by 4 meters. "The ball is red," you represent an object as red by placing it within a colour space. The yardstick performs a function, within such a system, that is in some ways similar to the function performed by a linguistic predicate. We represent something as standing in a relation to a given point on a scale (If you have a space where red is featured amongst all the colours, and place a red object in that location, you have categorized the object as red). And we could represent a different thing, b, as having the same length—same point on the scale; hence what a and b have in common (same type). But the scale is an unusual kind of semantic device, quite unlike a predicate term in certain respects. In particular, when you say something is red, you don't just say it's red, you say something about its relationship to the whole. As Wittgenstein describes this,

> just as all the graduation marks are on one rod, the propositions corresponding to the graduation marks similarly belong together, and we can't measure with one of them without simultaneously measuring with all the others.—It isn't a proposition which I put against reality as a yardstick, it's a system of propositions (1975: 110)

"The propositions joined by 'and' are not independent of one another, they form one picture and can be tested for their compatibility or incompatibility (ibid., 111-112). The concept of independent coordinates of description: the propositions joined, e.g., by 'and'

are not independent of one another, they form one picture and can be tested for their

compatibility or incompatibility. In my old conception of an elementary proposition there

was no determination of the value of a co-ordinate; although my remark that a coloured

body is in a colour-space, etc., should have put me straight onto this. A coordinate of reality

may only be determined once. If I wanted to represent the general standpoint I would say:

'You should not say now one thing and now another about the same matter.' Where the

matter in question would be the coordinate to which I can give one value and no more. The

general problem continuous quantities pose for sentential predication. The constituents of

sentences are *too atomic*.[28] "One shade of color cannot simultaneously have two different

degrees of brightness or redness, a tone not two different strengths, etc." "In which case,

propositions turn out to be even more like yardsticks than I previously believed.—The fact

that one measurement is right automatically excludes all others" (ibid., 110). The yardstick

presents a solution by offering a syntactic device which has no counterpart for the case of

linguistic syntax: we have a distinct representation of a 'property space,' but this is not

something whose existence (or non-existence) is asserted in the way that a singular and

predicate terms refer to objects and properties. Instead, the yardstick constitutes a third

category of thing: *form* or a *system* or a *framework*.

---

[28] According to Wittgenstein, when we try to capture statements of a degree of a quality using the model of sentential predication this leads to absurdities. He suggests an 'enumerative analysis,' by which a statement of degree is analyzed as a logical product of single statements of quantity and a completing supplementary statement—"(a)s I could describe the contents of my pocket by saying, "It contains a penny, a shilling, two keys, and nothing else" (1929: 167). For example, if b is the unit of brightness, then $E(b)$ is the statement that the entity E possesses this brightness. But then the proposition $E(2b)$, which says that E has two degrees of brightness, should be analyzable into the logical product E(b) & E(b)—which is just equal to E(b). Wittgenstein considers a way of answering this objection by distinguishing the individual units of the scale by assigning a distinct identity to each unit; for example, E(2b) would be E(b') & E(b''). This proposal manages to resolve the foregoing problem. But it also brings a new problem in its wake: if we assume two different units of brightness, then given an entity which possesses one unit, the question could arise, which of the two—b' or b''—it is; which is obviously absurd" (ibid., 168).

We may clarify Wittgenstein's proposal concerning the form of propositions about magnitudes by considering an analogous claim that has been made concerning the structure of magnitudes themselves. As a general characterization, magnitudes are features of the world are perspicuously represented using, or measured by means of, a numerical scale: amounts of mass, wavelengths of light, and spatiotemporal distances. According to Chris Swoyer, we should view measurement

> as involving properties of the sort that W. E. Johnson called determinates, specific lengths, like the property of being two meters long, rest masses, and so on. We also need higher-order properties and relations that can be exemplified by such determinate properties, for example the relation being longer than that holds between two determinate lengths just in case the first is longer than the second. This is not intended in any way to deny that our measurements involve individual objects, but only to assert that the facts discovered in such measurements involve properties. (1988: 243)

The view Swoyer is articulating in this passage is a response to a certain problem that arises when we attempt to characterize magnitudes themselves. We tend to think of the length of an object as an intrinsic property of that object. For instance, an object's having the length that is measured by 2 inches is a property that object has regardless of what is true about other objects. We also think of magnitudes like length as generating relations amongst objects. They generate an *ordering* on objects: a two inches-long object is longer

than a one inch-long object, and shorter than a four inches-long object. We can also say how much longer one thing is than another: a two inches long object is closer in length to a one inch-long object than it is to a four inches-long object. So magnitudes also generate *distance* relations between objects. Moreover, it seems as though these relations are ones that hold necessarily of any objects which have these properties: a two inches-long object couldn't fail to be longer than a one inch-long object or shorter than a four inches-long object. The problem is to account for the nature of magnitudes such that they essentially give rise to these ordering and distance relations amongst the objects which have them? One kind of answer that has been proposed is a relationalist account of magnitudes: magnitudes are nothing over and above the relations that hold amongst the objects that have them (Bigelow and Pargetter, Dasgupta, 2013). For example, an object's having the length that is measured by four inches is nothing over and above its being *longer than*, and *longer than by a certain amount*, certain other objects (a five inches long object, a 10 inches long object, and so on); and shorter than other objects (a two inches long object, a one inch-long object, etc). But the distinctive role magnitudes play in scientific explanation, which pose a problem for the relationalist.[29] Swoyer's proposal amounts to a realist view of magnitudes:

---

[29] Whichever account of relations we choose must, at some point, make indispensable reference to objects. Relations must have their places filled by terms, and the terms must at some point be non-relational, if we are to avoid an infinite regress. Since the relationalist denies the existence of properties, these can't fill the places; the only other choice is objects (a commitment which is reinforced by the relationalist's being motivated by considerations about observability). And the fact that the relationalist is, at some point down the line, forced to make essential reference, in its account of the properties of a given object, to facts about *other* objects, makes a relational account unsuitable for fulfilling the *explanatory role* that magnitudes are supposed to discharge. The main reason is that in many cases—most, even - the relation of a given object to other objects plays no role in explaining the features of the objects we are interested in. Consider the idea that the length of the shadow cast by a flagpole is explained by the length of the flagpole. The anti-realist is committed to the idea that the length of the shadow is partially explained by the objects which are the terms of this relation. Yet this seems absurd. As Peacocke puts it, "The properties of the standard gram or meter in Paris have nothing to do with, are irrelevant to, the explanation of why the avalanche flattened the forest or why the flagpole cast

lengths, masses, wavelengths of light and so forth constitute an ontologically basic category

of entities, as distinct from ordinary properties (Mundy, 1987; Peacocke, 2014). The realist

rejects the sharp distinction that is generally supposed to obtain between the functional

roles of particulars and properties, on which particulars are bearers of properties (whether

monadic or relational) and properties are instantiated by objects, but properties are never

themselves bearers of properties or relations. Magnitudes, on this view, are both an

intrinsic, monadic quantitative property and also fundamentally relational.[30] The second-

order relations hold necessarily. For example, if two grams mass is less than three grams

mass, then it is necessarily so; there is no possible world where three grams mass is less

than two grams mass instead.

Whether or not one accepts these claims about magnitudes, we can use them as a

basis for understanding how the model of a yardstick may provide an alternative model for

classification. A common way of characterizing representations of magnitudes (particularly

spatial magnitudes) is in terms of a relational model which is a kind of representational

analog to the position of the relationalist about magnitudes. This kind of relationalist

---

a certain length of shadow. This is reflected in the counterfactuals supported by such explanations. Other things equal, if the standard gram in the vault in Paris had been filed down, however much, the avalanche whose momentum had a certain magnitude would still have flattened the forest. (Peacocke 2014) In a similar vein, Brent Mundy points out that "Surely the whole system of physical quantities and quantitative laws would not collapse if through some cosmic accident all of the actual examples of objects precisely two meters long were to be destroyed while the standard meter itself remained intact."(Mundy 1987).

[30] To a rough approximation, the technical details of the proposal are as follows. We begin with the array of intrinsic, monadic quantitative properties. To these, we add two second-order relations: less than or equal to ($\lesssim$) and sum of (*). $\lesssim$ generates an ordering over properties; e.g., one gram mass is less than two grams mass, which is less than three grams mass, and so on. And * corresponds to a notion of "summation" over properties; e.g., the sum of one gram mass and two grams mass is three grams mass. Intuitively, this gives us the distance or closeness structure of mass quantities—distance between one gram and three grams is 2 grams. The axioms governing $\lesssim$ and * guarantee that these relations behave in appropriate ways. For details, see Mundy (1987: 37—40).

proposal is often found in representations of space, including so-called 'cognitive maps' (O'Keefe & Nadal, 1978). For example,

> Suppose, for example, that my cognitive map of Boston proves faulty: [...] my map had previously located an object *o* at location *l*. Now here I am at *l*, and *o* is not to be found; instead, u is there. What to do? It is clear enough what to say if we restrict our attention to the construction of a representation: I should remove the 'marker' that indicates o from its position on the map and put a 'marker' representing u there; I can then either put the o-marker somewhere else on the map or just leave it off. ... The relation between the contents of my maps before and after this change cannot naturally be described in terms of... structured propositions. ... (T)he problem is that moving the o-marker, for example, does not simply change where o is represented as located; o was also located in relation to other objects, and many of those relations—though not necessarily all of them—will have changed as well. It is thus not as simple as swapping one conjunct for another: The sorts of changes involved will be on a much larger scale. (Heck, 2007: 13)

The kind of holism illustrated by this example consists in the fact that the monadic property which an object is represented as having is determined by that object's relations to other represented objects. What it is for an object to have a given location just is for it to stand in a totality of relations to all of the other objects within that space. As Evans puts it, on a cognitive map, the simultaneous representation of the spatial relations of the objects

'constitute the frame of reference' (1982: 131). The tell-tale sign of this form of holism, as Heck observes, is that if you change or remove the features of certain parts of the representation, this will result in global changes to the significance of other parts.[31] What Wittgenstein is proposing, through his Yardstick theory of propositions, is a view of the contents of statements about magnitudes *which parallels the proposal of the realist:* we represent magnitudes as being both intrinsic properties of the objects and ones that bear constitutive relations to other properties of the same kind. The syntactic device of a scale locates the holism holds at the level of *properties*, rather than at the level of the objects which are the bearers of properties. The units of the scale are interrelationally identified; but the property we represent an object as having does not depend on the relations in which other objects stand to the scale. Hence, substituting one object term for another does not change the meaning of the attributive; and nor does (say) eliminating an object from a complex representation of this form. This difference in turn explains why the yardstick model is, whilst a standard relational spatial framework is not, suited to be a form of classification. Because, in a fully relational space, there is strictly speaking no such thing as the representation of the location of an object in isolation from all of the other objects and their features, such a relational structure lacks the kind of structure which would allow for the representation of different things as being of *the same type*. There is no part of my cognitive map, for instance, that could be common to a different cognitive map containing different objects; because the very existence of these differences changes the meaning of the elements within it.

---

[31] (The holism found in pictures is similar but more global still is often also attributed to pictures, or to 'iconic' representations more generally (Fodor, 2008))

We can draw an interesting parallel between the holism inherent in a scale and the kind of holism that is often attributed to predicative forms in thought and language. The 'context principle,' often attributed to Frege, says that "only in the context of a sentence does a word have a meaning."[32] The same idea has been extended to conceptually structured thought-contents:

The language case is useful also for illustrating this point: each of the abilities involved in the thought that a is F, though they are separable, can be exercised only in a (whole) thought and hence always together with some other conceptual ability. This is the analogue of the fact that the understanding of a word is manifested only in the understanding of sentences, and hence always together with the understanding of other words" (Evans, 1982: 102).

we should [reject] the idea that the contents one puts together in discursive activity are self-standing building-blocks, separately thinkable elements in the contents of claims or judgments. One can think the significance of, say, a predicative expression only in the context of a thought in which that content occurs predicatively. (McDowell, 2008: 7)

The suggestion here is the denial that concepts can occur independently of judgement; hence, that judgement can be adequately analyzed as the putting together of independently

---

[32] Although see Jeffrey Pelletier's (2001) paper 'Did Frege Believe Frege's Principle' , for a fascinating discussion of whether this was, in fact, Frege's position.

significant elements. Crucially, however, this should not be equated with denying either that sentences/propositions have constituents; or that these are 'atomic' in the sense described earlier.[33] The holism that is being asserted for predicative form is *not* the holism of a cognitive map, where the semantic dependency of the parts on the whole ensures that we cannot speak of 'constituents.' Rather, the claim is a sentence / proposition is what we might call a 'unity': a genuinely complex object (so that we can meaningfully speak of a judgement as being *composed* of concepts, and of two distinct judgements as containing *the same* concept), but in which the whole is nevertheless somehow prior to or more fundamental than the parts; a whole whose constituents stand in internal or necessary relations, rather than external or contingent relations. A predicate concept is such that it both does and does not have a determinate meaning of its own: it has a meaning of its own inasmuch as it can be recombined with other constituents to generate novel propositions to which it makes *the same* semantic contribution; but it does not have a meaning of its own in the sense that it makes a semantic contribution when it is combined with a singular concept *or not at all*. The subject and predicate terms of a sentence or proposition are not, as McDowell has put it, 'self-standing building blocks' from which the whole is is constituted; these parts belong necessarily to one another. Now we can also make a similar observation about the syntactic device of a scale: a scale is likewise a 'unity' in the foregoing sense: it is a complex object whose constituents necessarily belong to one another; and whose relations are *internal*. A scale is not simply a collection of objects, in the sense that we cannot think of the units of a scale in terms of individual building blocks that

---

[33] There is a worry lurking here about how semantic compositionality and the context principle can be true together.

make an independent contribution to the scale, and can, accordingly, be independently satisfied. Rather, as Jose Medina has put it, the scale "is a relational structure whose elements are essentially interdependent. The units on a scale can only be identified as relative positions within the scale; that is, the identity of each unit presupposes *the whole scale*" (Medina, 2002: 34-35). Moreover, according to Wittgenstein, the considerations which motivate a recognition of this form of unity in the representational contents of statements about magnitudes are analogous to those that motivate propositional unity. The form of a proposition should be able to capture the difference between false propositions and nonsense.

Although I think the syntactic device of a scale represents a promising avenue for exploring alternative forms of classification, there are legitimate worries about the coherence of this proposal in the form Wittgenstein suggests, that is, as a proposal about *propositional* form. One problem can be appreciated by remarking on a further difference between the yardstick and a sentence: these differ not only with respect to their 'general elements,' but also with respect to their 'object terms.' The sentence, 'that stick is seven meters long' applies a predicate term, 'is seven meters long,' which denotes the property of being seven meters long, to an object term, 'that stick,' whose denotation is a particular stick. The term 'is seven meters long' is not identical with the property, being seven meters long, which it represents. And no more is the term, 'that stick' identical with the object which *it* represents. Now consider the yardstick. The analog for the predicate term, 'is seven meters long' is, as we have seen, a particular location on the scale: namely, the point on the scale which denotes the length *is seven meters long*. What is the analog for the

singular term? It is the stick itself. It by virtue of the fact that the stick has a particular length—hence occupies a particular extension within the scale—that the stick is 'represented as' (or, as we might say, measured as) having the property of being seven meters long. As an illustration of this difference, notice that in the sentence, the singular term 'that stick' does not constrain a particular predicate. A sentence could represent any size in relation to it, regardless of the nature of the stick. But in the present system, it does. Given the size of the object itself, and the nature of the scale, it could *only* be represented as having that value. Anything which has a size will have an interpretation relative to that system. And only things which have a size will have an interpretation relative to that system. What follows from all of this is that, although we can distinguish between the location on a scale and the property that it denotes, there is no corresponding distinction to be drawn between the 'singular term' of the yardstick system and the object that is so represented. The object itself *is the singular term.* In the light of these observations, it may occur to someone to point out that what we have in the case of the yardstick is so different from a sentence that it is questionable whether we can really speak of representation in this case. What we have is rather a system of measurement: the object's properties are *measured by* the object's relation to the scale.

A second problem lies in the suggestion that this kind of structure can yield a kind of *structured proposition*. We have seen that we must not confuse the representational content of a state, an abstract object which may be shared in common by different thoughts, had on different occasions or by different people, for the physically instantiated state with a spatio-temporal location. When we say that propositions have a sentence-like

structure, we are drawing an analogy between the mereological structures of sentences and those of abstract objects: that two different kinds of things, one physical, the other abstract, share a common part-whole structure is perfectly comprehensible in this case. The problem with the present suggestion is that the distinguishing features of predication in graphs in partly rooted in facts about the structures of their vehicles of content; that is, how the *signs* signify. In particular, the fact that graphs, as we have seen, refer to their objects by virtue of *shared qualities* of the signs themselves (magnitude structure of spatial dimensions). The distinctive features of graph-like predication derive from this fact. For this reason, it can seem hard even to make sense of what it would mean for a content to have a 'graph-like' structure. If the syntactic structure of a graph is inseparable from the spatial properties of those representations, how could this kind of structure possibly have a mereological analog for the contents of perception—which are by definition *non* spatial? In short, the worry is that we are at risk of making a 'vehicle-content' confusion.

To summarize, we have seen that the image of a yardstick provides a promising basis for understanding how there could be a kind of attributive structure *that is not a form of predication*. The yardstick model is non-atomic in the sense that it involves a space which includes information about how a property stands in relation to the whole system. Given that we have defined conceptual classification in terms of predication, the present proposal amounts to a kind of nonconceptual classification. But although representations which use the device of a scale form a distinctive and interesting subclass of *representations*, we have not shown that they demonstrate an alternative form of classification. Conceptual classification is a semantic notion, rather than a syntactic one. Very different kinds of

representations (a picture and a sentence, for example) can in principle have the same kind of content. In order to fulfil this promise, however, we need some way of understanding how this proposal can be applied at the level of representational content.

2.5 The Argument from the Nature of Propositions

The discussion of the previous section forms a natural basis for considering a second line of argument in favor of monism. This argument seeks to establish a constitutive connection between attribution and predicative structure based on the nature of propositions.

One way to introduce this argument is by revisiting Charles Travis's discussion of 'Frege's Line,' which we considered in the previous chapter. Travis's commitment to monism is demonstrated in the way he defines the notion of a concept:

> The key feature of the conceptual… is that for anything conceptual there is a … range that is the range of cases, or circumstances, which would be ones of something instancing that generality (or, again, a range of things not instancing it. (2007: 231)

For Travis, a concept is the semantic analog of a property: just as a property (being red meat) is an abstract object which can be instantiated by a range of different particulars, so a concept (the concept of being red meat) is an abstract object which can 'fit' (be asserted of or predicated of) *a range* of distinct particulars, in the sense that what 'fits' the concept red meat "might still have done had it had a bit more gristle, or, while older, were not overly oxidized, and even if Texans were all vegetarians." For Travis, then, anything which

is an attributive is a concept. Travis further holds that concepts are necessarily elements of structured propositions. For concepts (like *being red meat*) belong, together with propositions (like, *that there is a piece of meat on the rug*), on the left side of 'Frege's Line.' And, in belonging on this side of the line, they stand opposed to sensible particulars (like *that piece of meat*), which instance conceptual generalities and which are truth-makers for propositions. That something is an object capable of instantiating concepts entails that it is not the sort of thing that could be a concept ("A piece of meat is not in the business of being instanced. So treating it would be bad grammar.") Likewise, that something is a concept entails that it is not the kind of object that could instantiate concepts. In short, Travis's framework commits him to a substantive view of monism on which any entity capable of performing the semantic function of classifying must be a fully abstract constituent of a propositional content.

I believe we can understand this conclusion as deriving the fundamentally conceptual nature of classification from a prior commitment to the fundamentally propositional nature of representation. Begin with the idea that representational contents are *propositions*: non-sensible, timeless abstract entities individuated by a truth-condition. It is in virtue of this fact that a proposition can be common to many different states: a sentence and a thought can have the same content; you and I can think about the same thing; I can think about the same thing on different occasions.[34] As noted in the previous chapter, this idea that representational contents must be mind-independent (i.e., non-

---

[34] Notoriously expressed through his platonism, in which propositions inhabit a 'third realm' of causally inert objects, between the world of actual objects and the world of subjective mental states. (See 'The Thought'). But we don't have to think of propositions in this way - as causally inert and magically correlated with a physical world with which they are magically correlated. [Burge, Peacocke].

subjective) *sharable* entities is one of the insights with which Frege is most closely associated, and the basis of his attack on psychologistic accounts of mental representation. This observation provides a context for understanding the 'atomism' of concepts: their context-independent character is the sub-propositional counterpart of the context-free character of whole propositions. On Frege's view, as Tyler Burge points out,

> the truth-value of the cognitive content—of what the thinker grasps, thinks, or believes—is eternally fixed given its nature, given the kind of content it is. A trademark of a sense or Fregean thought component is that it can in principle be expressed on indefinitely many occasions. For nothing in its expression or in its being thought affects its referential relations. Its relation to its referent(s) is atemporal and depends purely on its own nature and the inventory of the world" (Burge, 1979: 238).

We can marshal these observations into an argument for why all classification must be conceptual in the following way. First, any mental state which functions to classify an object must be a relation to a proposition (because classification is a representational function, and representation is a matter of bearing a relation to a proposition). Second, the proposition must be complex, being composed of constituents which refer, separately, to objects and properties (this is a requirement because there must be types which can be shared in common by different propositions). Third, any constituent of a proposition must refer in a context-independent way (because otherwise it could not be the constituent of a

proposition; a proposition could not lack a spatiotemporal location if its constituents did not). In other words, the essential connection between predication and attribution reflects the idea that what is true for propositional wholes must also be true for their constituents at the level of whole contents. The essentially 'word-like' structure of attribution reflects the fact that all representational contents are abstract objects.

2.6 Attribution without Propositions

In section 2.3, we saw that in making his 'Yardstick Theory' a theory of *propositions*, Wittgenstein attempts to fit the representational device of a scale into a model of content where it does not naturally belong. The clue to this was the observation about 'measurement'; the role played by objects themselves in the model of a scale has no counterpart in the theory of propositions. In spite of these worries, I believe that Wittgenstein was fundamentally on the right track. His mistake was to overlook the full implications of the representational structure he identified. This is not a reflection on the inadequacy of the theory, so much as the limitations of the propositional model of representational content. The 'yardstick theory' is not an alternative view of propositional structure, but an alternative view of the structure of representation; one that is both abstract and particular at the same time. Wittgenstein could have benefitted from paying closer attention to Kant. Kant's theory of intuitional content gives us a different way of thinking about how a scale-like structure may figure in representational content, which I believe demonstrates the way in which such structures may give rise to non-propositional

forms of representation. By appreciating these ideas, we will see our way to a pluralist view of classification.

Kant recognizes two basic and irreducible kinds of representations of objects: concepts, which are non-sensory representations of universals from which judgements are composed; and intuitions, which are sensory representations. His distinction is based in part on the structures of these two forms of representation. A 'discursive concept' is an indirect or general representation of objects, which classifies these by bringing or 'containing' these '*under* itself.' Moreover, Kant follows Locke in regarding discursive concepts as formed through a process of, first, abstracting and, then, combining with one another the features that a number of objects are observed to share in common. In other words, a concept has the structure we have been associating with an atomic predicate-term. On the other hand, an intuition is a direct representation of objects, which is composed of a purely subjective, sensational element ('the effect of an object on the representative capacity, so far as we are affected by it' (A19-2O/B34); and a 'pure' or abstract 'form,' which in the case of outer perception, is *space*. And Kant explicitly distinguishes the structure of space from that of a discursive concept through its part-whole structure. In the first place, whereas a discursive concept contains an infinity of things under itself but is itself essentially simple, space "contains an infinity within itself." Moreover the parts of space "are possible only in the whole, not the whole through its parts": space is "essentially one; the manifold *in it*,' and we represent parts of space (say, a location, or region, or the extension of an object) only through 'bringing in limitations' or as

'determinations' of this whole.[35] Since intuitions are representations of objects which are partly composed of such a space, together with 'the real of sensation,' it follows that intuitions have a 'nonconceptual structure.' In other words, Kant distinguishes perception from conception on the basis of a contrast of precisely the sort we have been considering.

What is interesting about Kant's account, for present purposes, concerns certain conclusions he goes on to draw, on the basis of these differences, between the contents of concepts and intuitions. In his commentary on the transcendental deduction of space, Kant makes a point of distinguishing intuitions, *qua* sensory representation, from concepts, *qua* intellectual (non-sensory) representation, based on how they present their objects. A discursive concept of (say) rightness, presents that property as 'a thing in itself.' Kant contrasts the contents of concepts in this regard with intuitions: which represent *appearances*; that is, subjective representations of objects. As Kant explains this point:

> this presentation contains nothing whatever that could belong to an object in itself. It contains, rather, merely the appearance of something, and the way we are affected by that something. This receptivity of our cognitive capacity is called sensibility; and even if we were to see through that appearance and to its very bottom, yet this receptivity remains as different as day and night from cognition of the object in itself.
>
> ...

---

[35] Indeed, Kant explicitly identifies space, alongside judgements, as an example of a 'unity': 'Space should properly be called not compositum but totum, since its parts are possible only in the whole, not the whole through its parts' (A438/B466).

Hence sensibility does not merely fail to provide us with a distinct cognition of the

character of things in themselves; it provides us with none whatsoever. And once we

remove our subjective character, then the pre-sented object, along with the

properties contributed to it by sensible intu-ition, is not to be found anywhere at all;

nor can it possibly be found, be-cause this subjective character is precisely what

determines the form of that object as appearance. (B62/A45)

I believe we can understand this notion of 'the thing in itself,' which Kant uses to describe

intellectual or conceptual content, as the analog of a Fregean sense: it is a context-

independent, or fully conceptualized grasp of a content. To use an example Kant gives, the

concept of 'rightness' refers to that feature in a *fully objective* way, independently of how it

affects us. When Kant claims that intuitions present us with 'mere appearances' and not the

'object in itself,' he is not saying that intuitions lack an objective content, however. For an

appearance, unlike a sensation, an appearance *of something*. Rather, he is saying that the

content of intuition, although objective, is *also subjective*: it gives us a presentation of an

object *through its effect on us*. To put the point another way, the space / sensation dualism

applies not only at the level of intuitions, or perceptual *states,* but also at the level of

appearances, or perceptual representational content. An appearance is a kind of

representational content that is partly *constituted* by a sensation. Kant makes this point

explicitly in the Transcendental Logic: "*Appearances, as objects of perception*, are not pure

(i.e., merely formal) intuitions, as space and time are (for these cannot in themselves be

perceived at all)"; rather "appearances contain, in addition to [pure] intuition, the matter

(through which something existent is presented in space or time) for some object as such"

(B208). In short, when Kant characterizes intuition as presenting us with merely an

appearance of an object, he is departing from the Fregean line considered in the previous

chapter, on which propositions exist in a 'third realm' that is sharply separated from the

'subjective realm.' He is ascribing to intuition a *non-propositional* (hence nonconceptual)

representational content: one containing, together with an abstract component; a concrete,

sensory component; "the real of sensation." An appearance is a representational content

that *contains* 'the real of sensation': "Besides intuition, appearances, as objects of

perception, contain the material for some object in general, that is, the real of sensation,

through which one can only become conscious that the subject is affected."[36] As distinct

from a mere sensation which pertains only to the subject. Instead, I believe he is saying that

because of their structure, intuitions refer to objects in a fundamentally subjective way. But

the nature of the representational vehicle guarantees that this condition will not be met.

For there is no way, within such a system, to pick out an attribute independently of how

that thing affects us. Our ability to refer to that property, within such a system of

representation, essentially depends on a sensation's occupying a certain space within the

frame of reference. As such, a perception is not simply a combination of a sensation and a

---

[36] Kant also draws a connection not only between *intuitions* and magnitudes, but also between *appearances* and magnitudes:

> Appearances are one and all magnitudes—specifically, extensive magnitudes, because as intuitions in space or time they must be presented through the same synthesis whereby space and time as such are determined.
> Extensive is what I call a magnitude wherein the presentation of the parts makes possible (and hence necessarily precedes) the presentation of the whole. I can present no line, no matter how small, without drawing it in thought, i.e. without producing from one point onward all the parts little by little and thereby tracing this intuition in the first place. (A160/B205)

separate representational content, individuated independently of the sensation. Rather, the content of an intuition is such that it is *inseparable* from the subjective component *because it is in some sense partly constituted by it*. It is not the case that there is a sensation and a separate representational content, which would still exist even if the subjective element were removed.[37]

On what basis does Kant infer, from the structure of perceptual representations (intuitions), these conclusions about perceptual contents (appearances)? I believe we can answer this question by reflecting on the structure of the perceptual problem itself. In a 1953 paper for the American Scientific Association, Warren Weaver surveys the progress of scientific thought. To this end, he distinguishes two kinds of problems each presenting intellectual challenges that differ in their level of difficulty: 'problems of simplicity' and 'problems of organized complexity.' Whereas problems of complexity involve a significant number of variables which are all interrelated in lawful or organized ways, and are consequently rather difficult, problems of simplicity are two-variable problems in which one quantity (say, a gas pressure) depends primarily upon a a second quantity (say, the volume of the gas); hence in which the behavior of the first quantity can be described, with a useful degree of accuracy, by taking into account only its dependence on the second

---

[37] Compare this notion of space as "subjective" with Gareth Evans's description of allocentric space as objective: "To say that the fundamental level of thought about the spatio-temporal world—the level of thought to which all our other thinking directs us is thought which would be sustained by a cognitive map of that world is to stress that our fundamental level of thinking is, in a certain sense, `objective.' Each place is represented in the same way as every other; we are not forced, in expressing such thinking, to introduce any 'here' or `there.'" (It is often said that in such thinking we are taking the third-person, or God's-eye, point of view, but … I reject this way of looking at the matter. … (I)n fact the thinking is truly objective—it is from no point of view.)" (1982:

quantity, and by neglecting other minor influences.[38] One might think of the problem perceptual systems encounter with regards to the representation of spatial features as being a problem of simplicity, in the following sense. The geometric properties of registrations of light on the retina are the joint product of features of the environment and features of a perceiver's perspective on the environment: spatial properties such as size, orientation, and location which objects have independent of being perceived; and relational properties such as distance from the observer, angle of observation, which are viewer-dependent properties of the object. One consequence of this fact is that any change in perspective will result in a corresponding change in the geometric properties of the retinal image, despite the properties of the object remaining stable. Spatial constancies are capacities to extract (say) an object's size from a signal which confounds this with viewing distance. This is a matter of working out how the stimulus decomposes into different dimensions. As Kant recognized, this problem is in some ways less simple than it seems. In order to be able to resolve it, one needs some internalized conception of space in order to be capable of extracting, from sensory registrations, spatial properties. One cannot read them directly off the stimulus. There must be something supplied by the mind: a 'frame of reference.' But it is also noticeably different from and easier than problems of complexity with which we engage in inferential thought. Although there is no straightforward correlation between a given size or shape on the retina and a given size or shape in nature,

---

[38]A grasp of problems of simplicity allowed for the development of scientific theories of light, sound, heat and electricity in the 17th and 18th centuries. But a range of scientific problems, particularly those in the life sciences such as biology and medicine, "could not be probed by this analysis." For example: "What makes an evening primrose open when it does? Why does saltwater fail to satisfy thirst? … What is a gene, and how does the original genetic description express itself in the developed characteristics of an adult? These are all problems which involve dealing simultaneously with a sizeable number of factors which are related into an organic whole" (1953: 49).

there is a correlation amongst the *kinds* of properties that are indicated. The size of a component of the retina may be a composition of indefinitely many different combinations of the size of an object and its distance from the subject, but it is always the joint product of *some size* and *some distance*. The problem of perception is the problem of decomposing the proximal stimulus in the right way; figuring out what goes where.

We can draw together this line of reflection with the theory of nonconceptual content. Perception is not determined by the stimulus, but it is *constrained* by the stimulus; inasmuch as perception makes essential use of sensory magnitudes in order to determine magnitudes in nature. Perception is a context-dependent representation of particulars and properties: a partial abstraction which is nevertheless dependent on concrete features of the context of perceiving. It is possible for perception to represent spatial properties in the distinctive way Kant suggests—by means of an innate, unitary spatial frame of reference— because the sensory registrations which give rise to spatial perceptions are complex magnitudes containing all the variables to be figured out within that system. The problem faced by the perceptual system is more complex than that involved in mere information registration, because the precise values along each dimension are underdetermined, and have to be figured out through internal processing. But it is less complex than those involved in reasoning, because the stimulus contains two factors that are directly related to one another in their behavior (say, size and distance; orientation and shape). The behavior of the first quantity can be described, with a useful degree of accuracy, by taking into account only its dependence on the second quantity. Perceptual systems work by exploiting these constraints present in the stimulus. This results in a representation which is abstract,

but *context-dependent*. It results a nonconceptual content: an abstract representation of a property can be combined with, and essentially dependent on, the particular which it classifies; in the sense that it depends on that object 'taking up' a certain amount of space.

In a certain sense, then, perception conceived in this way *is* a bit like measurement. That is not to say there is no interpretation going on. But it is the magnitude registered that is compared with a frame of reference; it does not involve the generation of a separate, independent meaning as happens with concepts. As Wittgenstein put it in connection with a different issue, "a photograph is not like a blueprint."

Let us take stock of the foregoing discussion, in relation to the broader argument of this chapter. The previous section introduces a suggestive picture of 'spatial attribution' to counter the argument from mereology: although discarding or leaving out information is one form in which classification occurs, we can also have classification in a form which preserves specificity. This section has extended this framework to counter the argument from the nature of propositions: although some forms of attribution occur within propositions, there are non-propositional forms of attribution as well. The yardstick model considered in relation to perception falsifies the propositional model, because it shows how an abstract representation of a property can be combined with, and essentially dependent on, the particular which it classifies. Taken in conjunction with one another, these two lines of argument give rise to a pluralist account of classification, on which we can recognize, along with *predicates*, a different kind of attributive in the form of a *property space*, or *location in property space*.

2.7 Comparing Conceptual and Nonconceptual Classification

Up to this point I have focussed on the monist as someone who is committed to an essential connection between attribution and atomism. Against such a view, I have argued that there are two possible forms that attribution may take: predication, which *brings objects under* an atomic representations of a type; and 'spatial attribution', which *places objects in* a property space or frame of reference, at a location which refers to that type in a fundamentally holistic way. Yet there are some philosophers who operate with a notion of 'concept' that is broader than the atomistic picture I've been focussing on. These philosophers may grant the distinction I've defended whilst denying that what we have is a deep difference between two kinds of content. Such a stance invites a deeper kind of reflection concerning the significance of distinguishing forms of classification. When all is said and done, why should we be interested in the difference between atomic and context-specific forms of classification?

John McDowell's more recent work exemplifies the kind of challenge I have in mind. Whilst McDowell has consistently maintained that perceptual experience must be conceptual, his view about what that would mean has shifted over the years. Circa *Mind and World*, a 'conceptual content' is defined in relation to judgement; and experiences have a conceptual content insofar as the contents to which one is related in experience are those very things that could be the content of a judgement. However, in his later work he articulates a more nuanced position: the feature we should care about with regards to concepts is *classification* regardless of how this function is realized; although there may be differences, any of these differences are not as important as the similarities. In 'Avoiding

The Given,' he acknowledges that judgement and experience have a different kind of structure. Judgement has a discursive structure ("we can think of judgments as inner analogues to assertions" (2008: 8)); and we derive our concept of a 'conceptual content' most centrally from this kind of 'discursive' content. Experience has an 'intuitional' rather than a 'discursive' structure: "the unity of intuitional content is given, not a result of our putting significances together. Even if discursive exploitation of some content given in an intuition does not require one to acquire a new discursive capacity, one needs to carve out that content from the intuition's unarticulated content before one can put it together with other bits of content in discursive activity. Intuiting does not do this carving out for one." Intuitional content, since it lacks a discursive structure, is not the sort of thing one could judge. But if this is right, why go on insisting that the contents of experience are conceptual?' Says McDowell,

> Because every aspect of the content of an intuition is present in a form in which it is already suitable to be the content associated with a discursive capacity, if it is not— at least not yet—actually so associated. That is part of the force of saying, with Kant, that what gives unity to intuitions is the same function that gives unity to judgments. If a subject does not already have a discursive capacity associated with some aspect of the content of an intuition of hers, all she needs to do, to acquire such a discursive capacity, is to isolate that aspect by equipping herself with a means to make that content—that very content—explicit in speech or judgment. The content of an intuition is such that its subject can analyse it into significances for discursive

capacities, whether or not this requires introducing new discursive capacities to be

associated with those significances. (2008: 9)

I understand the claim here to be that intuition, although somehow 'essentially whole,'

rather than made up of separable 'terms,' nonetheless has a 'logical form.' It has a subject-

property structure; the world is already classified into objects and properties. What

happens in the process of converting from intuitional to discursive content is the addition

of *atomism*. The shift in position marks a shift in McDowell's conception of *what is central*

to the notion of 'the conceptual sphere': although we should still "centre our idea of the

conceptual on the content of discursive activity," we should not think of discursive activity

as exhausting the conceptual. Rather, what matters is that 'what gives unity to intuitions is

the same function that gives unity to judgments.'

I think such a stance is worthy of reflection. In effect, it raises questions which often

go unasked in discussions of 'nonconceptual content': What is a concept, really? What are

the theoretical interests that determine our characterizing contents as being of certain

kinds? And that would license difference of *kind*—as opposed to variations within a kind?

These questions would not be answered by appealing to a theoretical account of the nature

of conceptual content in thought. For what is being asked, in effect, is for a more general

account of what the concept of 'concept' means for us, and why we care so much about

whether a representation *is* one, or whether there could be representations without such

things. To put it in different terms: suppose we accept that the conceptual / conceptual

distinction marks a duality in the mind, with two fundamentally distinct kinds of things on either side of it—then what reason do we have to regard this distinction as one like *that*?

As I noted in the introduction, we invoke the notion of a concept to explain certain capacities a subject has, most notably (although not exclusively) their capacities for reasoning. When we consider the ways in which a thinker's inferential capacities depend on conceptual structure, One thing that matters is their *logical form*: the semantic structure of content separates the overall content into objects and types of properties in a way that makes transparent the validity-preserving contributions these make to truth-values. Yet it is not only the fact that concepts group and type objects as being of some kind but the fact that they do so *atomistically* that matters for these explanatory purposes. Classification must be realized through atomic, recombinable conceptual constituents in order for subjects to be able to *make* inferences; to represent generalizations and to draw appropriate conclusions. Spatial forms of attribution, like concepts, represent objects as belonging to types. They also allow for the recognition of semantic structure—commonalities amongst a creatures mental states. Yet spatial forms of attribution, being non-atomistic or context-dependent, do not allow for these further capabilities provided for by atomic representations of categories. Since spatial attribution is compelled to represent magnitudes in conjunction with one another, this form of classification doesn't support generalizations. And since the predicates can't be freely detached from the objects which are presented as being that way, this form of classification doesn't support the drawing of conclusions in inferences. We may conclude from this that the atomic structure of conceptual classification makes possible certain capacities which would be closed to a

creature possessing only representations organized in the form of spatial attribution. A creature capable only of spatial attribution would not be capable of propositional inference.

Since we are discussing McDowell, it is worth noting how these comments bear on McDowell's arguments against nonconceptual content which we considered in the previous chapter. For McDowell, it is a condition on an adequate theory of perception that it be capable of explaining how experience can provide reasons for our beliefs about the external world. McDowell's commitment to a conceptualist view of experience follows from this 'metatheoretical' criterion about what a theory of perception must explain: he holds any mental state we locate outside the sphere of concepts will be incapable of supporting 'rational linkages.' But the foregoing observations cast doubt on this conclusion. Spatial attributives have a *logical form* which separates the overall content into objects and types of properties in a way that makes transparent the validity-preserving contributions these make to truth-values. Yet spatial attributives differ from the conceptual constituents of thoughts in ways that are relevant to explaining aspects of rational capacities apart from the appreciation of rational relations. McDowell's argument *implicitly relies on the monistic view of classification* inasmuch as it assumes that any representation which classifies the world as being some way will be a concept. Once one acknowledges the possibility of classification without concepts, we see that there is no inconsistency between holding a nonconceptualist view of experience and respecting the idea that experiences must be capable of standing in rational relations to judgements.

A further difference concerns what the kinds of properties that could be represented by means of these different systems. The atomic structure of conceptual

classification makes it possible for us to discover and represent aspects of the world that are not available to us within a system of spatial attribution. Hume's *Enquiry* describes the remarkable ability thinkers have to transcend the limitations of their embodied perspective: "What never was seen, or heard of, may yet be conceived; nor is anything beyond the power of thought, except what implies an absolute contradiction" (1975: 18). The transcendent nature of thought "(t)o form monsters, and join incongruous shapes and appearances" and to "transport us into the most distant regions of the universe; or even beyond the universe, into the unbounded chaos, where nature is supposed to lie in total confusion" is possible because the conceptual constituents of thoughts are capable of being detached from the contexts in which we learn to apply them. But spatial attribution is doubly context dependent, requiring both that a property be represented in relation to a system of properties and that it be represented in relation to an object which *is* that way. As such, the reach of spatial attribution is limited to properties like magnitudes which belong to a system and which are capable of being directly registered by a subject through their senses.

It would be a grave mistake to suppose that the gains are all on the side of thought, however. Nietzsche, whom I quoted earlier on the atomic nature of conceptual classification, uses the atomism of concepts to launch an excoriating critique of the distorting influence of language and conceptual thought. Although "we believe that we know something about the things themselves when we speak of trees, colors, snow, and flowers," he says,

we possess nothing but metaphors for things—metaphors which correspond in no way to the original entities. In the same way that the sound appears as a sand figure, so the mysterious X of the thing in itself first appears as a nerve stimulus, then as an image, and finally as a sound. Thus the genesis of language does not proceed logically in any case, and all the material within and with which the man of truth, the scientist, and the philosopher later work and build, if not derived from never-never land, is a least not derived from the essence of things. (1871: 119)

It's interesting to observe that whereas Kant calls the content of a concept 'the thing in itself,' Nietzsche's criticizes concepts precisely for *failing* to represent the "thing in itself." In spite of appearances, there is no real disagreement here. Nietzsche accepts Kant's idea that the concepts are atomistic presentations of properties, discarding contextual details for a content shorn of a subject's perspective on them. His point is that the 'thing in itself,' conceived in these terms, is in some ways a distortion of the concrete reality that is being thought about. Greater objectivity in representational content—the adoption of an increasingly abstract view of the structure of the world shorn of our perspective on it—does not always lead to something that is more 'true.' When he writes, immediately following on from this passage, that "the pure truth, apart from any of its consequences, … is likewise something quite incomprehensible to the creator of language and something not in the least worth striving for," he is highlighting deficiencies in how intellectualist or rationalist philosophies have conceived of the truth. In the context of the present discussion, what is interesting about this is that it makes space for us to recognize some of

the advantages that attach to spatial forms of attribution as compared with their conceptual counterparts. With conceptual abstraction, we get enhanced freedom and control over how we use our existing representations to generate new ones. Yet one might think that the trade-off here is a loss in our appreciation of the structure of a basic layer of reality. The syntactic constraints inherent in spatial attribution reflect a kind of detail and specificity that inheres in the very properties that are being represented. On the other side of this equation, then, the 'real of sensation' participates in non-propositional content marks an encoding of the concrete reality of the object of perception in a way that is lost when we transition into atomistic representations of singular terms. If spatial attribution is more limited, that is in some ways because it is more tightly connected with the concrete reality it represents.

David Marr has claimed that what is "fascinating and powerful" about the idea of representations as a means of 'describing' reality is the realization that how information is represented can greatly affect how easy it is to do different things with it. He illustrates this point with the example of different systems for representing numbers.

> if one chooses the Arabic numeral representation, it is easy to discover whether a number is a power of 10 but difficult to discover whether it is a power of 2. If one chooses the binary representation, the situation is reversed. Thus, there is a trade-off; any particular representation makes certain information explicit at the expense of information that is pushed into the background and may be quite hard to recover. ... It is easy to add, to subtract, and even to multiply if the Arabic or binary

representations are used, but it is not at all easy to do these things—especially

multiplication—with Roman numerals. This is a key reason why the Roman culture

failed to develop mathematics in the way the earlier Arabic cultures had. (1981: 42)

The answer I have been sketching to the question of why this difference in forms of

classification is theoretically meaningful similarly builds on the connection between the

form of a representational system and its implications for what subjects are able to do with

their representations. A conceptual representation of F and a representation of F by means

of a space both represent the property, is F. Both do so in a way that allows for a semantic

structure in a subject's representations; facilitating the recognition of sameness. But there

are many differences between them in terms of what *else* they allow for, beyond this. The

capacities which conceptual classification explains are different from the capacities which

spatial classification explains. They are different explanatory kinds.

2.8 Conclusion

Is classification is essentially conceptual? Are we entitled to treat as equivalent the notions

of 'representing x as being of the type, F' and 'bringing x under a concept of F'? This chapter

has given a negative answer to these questions. Concepts are essentially ways of classifying

objects, but classification is possible without concepts; that is, without atomic propositional

constituents. My argument for this conclusion proceeded in three stages. First, I showed (in

section 2.4) that structured representations composed of singular and general terms may

have an organization that is not 'sentence-like' but more 'space-like.' In a sentence, a

general term is a word or phrase, which refers to a given property atomically, without reference to anything that is not that property. In a coordinate space, a general term is a location in a 'property space' which encodes the relation that magnitude bears to a system of magnitudes to which it belongs. Next, I showed (in section 2.6) that 'space-like' structure can be realized at the level of abstract, representational contents to yield a non-propositional content; partly abstract, partly concrete and sensuous. Finally, I showed (in section 2.7) that these two forms mark a theoretically meaningful distinction in terms of what a subject is able to do with their representations of categories. Spatial attribution, which doesn't permit the kind of recombinability and stimulus-independence that we find in atomic concepts, is 'nonconceptual' in a theoretically meaningful sense. It precludes many of the capacities concepts are invoked to explain.

To claim that there *could* be a mental state which classified objects without using concepts is not yet to claim that perception, or any other kind of mental state, is this way in fact. In this regard, the conclusion of this chapter may seem rather modest. It may be viewed as an invitation to further investigation, rather than a 'last word' on this subject. Nonetheless, even to know that nonconceptual classification is possible turns out to have some rather significant consequences for the nonconceptualism debate.

1) The possibility of classification without concepts carries implications for how we should think of nonconceptual content. Typically, discussions of nonconceptual content have operated with a conception of 'nonconceptual' on which this is some kind of unstructured content: 'nonconceptual' content must be content which lacks an attributive structure, hence which does not involve classification. Nonconceptual classification

provides us with an alternative way to understand the idea of a nonconceptual content. This notion of nonconceptual content is, moreover, eminently suited to be the content of a perceptual experience. I believe it is preferable to think of nonconceptual content as a form of classification, since this conception easily copes with some of the worries that have dogged this more common alternative. As we saw in the previous chapter, many who oppose a nonconceptual content on the grounds of incoherence. Some object to the very idea of a nonconceptual content because they doubt that something could perform the function of representational content—to specify correctness conditions—without involving the representation of classes or types (Brewer, 2006; Travis, 2007). Others object to the idea of nonconceptual content, especially in relation to perceptual experience, on epistemic grounds. If we take perception to have a different kind of content from beliefs and judgements, it is unclear how perceptual states can justify conceptual states, since it is difficult to conceive how different kinds of content could be inferentially related (McDowell, 1994; Byrne, 2005; Duhau, 2009). Monism has been a tacit assumption in both of these a priori arguments against nonconceptual content. Once we disabuse ourselves of the mistaken idea that the nonconceptual must preclude abstract classification, these objections lose their force. We see that representational content can be abstract without rising to the level of conceptual abstraction. And we also see that perception can have an attributive structure—hence, can serve as a reason for belief—whilst remaining outside of the 'sphere of concepts' by virtue of the lesser degree of freedom and flexibility it affords.

2) One might think that the mere possibility that perception might involve a form of nonconceptual classification provides *prima facie* reason to think that perception is this

way in fact. This point relates to our discussion of the nonconceptualism debate in the previous chapter. There, we saw that the arguments that motivate conceptualists and naive realists to their respective views of experience both rely in large part on negative arguments against the possibility of nonconceptual content. McDowell's epistemological argument for conceptual content relies heavily on the shortcomings of the perceived alternative: a naive realist alternative of experience. Accusing Evans of 'fraudulence,' McDowell equates 'nonconceptual content' with this kind of view. Meanwhile, Travis's epistemic argument in favor of naive realism relies heavily on the failures of conceptualism, which he regards as the only coherent alternative since the idea of a 'nonconceptual content' "makes no sense." These two opposing lines of argument both rely heavily on an assumption of what our options are when it comes to choosing an account of perceptual objectivity: either experience is non-representationally direct or it is conceptual; there is no middle way. The possibility of nonconceptual content reveals this to be a false choice: spatial attribution, both abstract *and* concrete, is fitted to serve as a middle way between a view of experience as purely concrete and one on which it is conceptually abstract. This discovery in turn counts in favor of a nonconceptualist view of experience: first, by directly undercutting the motivations for rival views; and second by promising a way of breaking the stalemate between them. By explaining *both* what the naive realist wants to explain and what the conceptualist wants to explain, the promise of nonconceptual classification removes the seemingly impossible choice between them. Of course, these are only prima facie considerations: it is open to the conceptualist or naive realist to respond with further arguments.

3) The model of nonconceptual content which I have provided promises to provide clarity to the question of how we might establish the existence of nonconceptual content. A fully satisfying argument for that conclusion will provide a clear positive criterion for what makes representation nonconceptual and show that perception satisfies it. This proposal suggests that a key source of answers lies in studying the psychological capacities associated with the attributives a subject represents. Conceptual attribution, by virtue of its atomic structure, ascribes (say) redness to an object in a way that leaves out all determining information concerning the specific shade of red that object has or its relations to other properties. The atomic structure of concepts grounds certain predictions about what a representing subject will be able to do with their concepts.—predictions of the sort Gareth Evans famously captured in his 'Generality Constraint,' which says that a subject who can think of that object as red will be capable of thinking, for any singular concept x which that subject possesses, that x is red. Since spatial classification involves attributives with a context dependent structure, however, a representation of red within such a system would not entail this sort of limitless flexibility or recombinability. Rather, one would predict systematic limitations on the ways in which a subject can apply or recombine such a content, respecting the constraints imposed by its place in a system and its ties to the causal effects of the objects. Even if one questions whether this kind of evidence is decisive—one might hold that such systematic limitations are consistent with conceptual structure by appealing to constraints imposed from outside of content (modularity within a creature's cognitive architecture)—it represents a promising place from which to start.

Along with these positive implications, I must also mention some limitations of the present investigation. I have largely been concerned to answer the skeptic about nonconceptual classification. This answer has consisted in an argument by counterexample: I've shown that monism is false by showing that there is a kind of content which is attributive without being composed of atomistic constituents. Yet this argument may be unsatisfying to someone already sympathetic to the possibility of a pluralistic approach to classification, who might well accuse me of taking a position that is still too close to that of the monist. For whilst presenting myself as a champion of a pluralistic approach to classification, what I've really championed is a dualism about classification. Whilst nothing I have said explicitly rules out alternatives to the form of nonconceptual classification I've proposed, my claim (in section 2.6) that there is some essential connection between 'anti-atomism' and 'context-dependence' carries an implicit suggestion that you can't get nonconceptual classification other than through this non-propositional marriage of a sensuous matter and an abstract organizing framework within which this is placed. Yet not only does this go against a central moral of this chapter—that we not rule out a priori the existence of different semantic strategies to the same route—but it would also seem to place unreasonably draconian restrictions on the applicability of nonconceptual classification to perception-like states containing a sensorious element. For the pluralist about classification, these are difficult and important questions whose answers await further investigation. Such a proposal would have to explain how Another question concerns the realization of nonconceptual classification within the psychology of individuals.

In summary, this chapter has made important progress in addressing the question of the nature of perceptual objectivity: first, by showing that nonconceptual content is a coherent and indeed attractive alternative to conceptualist and naive realist views of experience; and second, by giving some indication of what would count as evidence for this view. The next two chapters will build on these foundations. They represent a turn from theoretical issues concerning what the nature of perceptual objectivity *could be* to epistemic issues concerning how its nature can be known.

3. RETHINKING THE GENERALITY CONSTRAINT

3.1 Introduction

I have argued that perception could have a nonconceptual content: contrary to what some philosophers have suggested, this way of thinking about perceptual objectivity *makes sense*. But how could we determine whether perception possesses such a nonconceptual content as a matter of fact? What sorts of phenomena require explanation in these terms?

For those seeking answers to these questions, the Generality Constraint forms the basis for a promising project. In its original articulation by Gareth Evans, the Generality Constraint says that:

> If a subject can be credited with the thought that a is F, then he must have the conceptual resources for entertaining the thought that a is G, for every property of being G of which he has a conception. (1982: 104)

Evans's proposal, it is widely agreed, provides us with a criterion for assessing conceptual content in thought: what it is to think about a state of affairs *conceptually* can be understood in terms of the abilities a subject demonstrates to freely recombine the constituents of contents in the generation of new thoughts. Yet this positive criterion conceptuality might also serve double-duty as a negative criterion for representation *without* concepts: if a subject's abilities to recombine the parts of their representations is evidence of conceptual content, then systematic failures of generality might constitute

prima facie evidence for nonconceptual content. The value of the Generality Constraint for the friend of nonconceptual content thus lies in its promise to provide both a theoretically meaningful context for understanding classifications of contents into kinds and a practical way of determining those classifications for particular cases. Yet someone engaged in this project must ensure that failures of generality are consistent with the idea that what we are identifying about is truly a kind of representational mental state. My aim in this chapter is to highlight the ways in which philosophers engaged in this project have failed to meet this challenge, and to suggest how the Generality Constraint should be understood so as to rectify these shortcomings.

The formulation Evans uses in his 'official' statement of the Generality Constraint articulates the antecedent condition as a subject's being *credited with a given thought*. Now one way to interpret this is to take Evans as describing a condition which attached to thought by virtue of their conceptual nature. This gives rise to what I'll call a 'restricted' interpretation of the scope of the Generality Constraint:

> A subject's ability to recombine the constituents of their representations, as specified by the Generality Constraint is a condition on their having contents of a certain ('conceptual') kind. Failure to satisfy that condition would entail that the representation lacked a conceptual content. But the constraint is silent on the question of whether a state has a representational content *tout court*; i.e. a content of a different ('nonconceptual') kind.

The restricted interpretation suggest that even if all thoughts are as a matter of fact conceptual, this coinstantiation is a contingent fact about thoughts—since a thought's conceptual structure is evinced by a property that is independent of their basic function of being *about* states of affairs in the world. One might think it is in virtue of the fact that the Generality Constraint is restricted in this way that this constraint allows for the possibility of the existence of states with nonconceptual content. It is *only* if there is no a priori connection between *thinking* about the world and a thought's possession of conceptual structure that it can be an open question whether there could be mental states which have a content yet which do not involve the possession of concepts. For suppose, on the contrary, that subject's possession of a conceptual content, as evinced by their satisfaction of the Generality Constraint, is a necessary condition for ascribing contentful thoughts to subjects. Then the conditions which would qualify a thought that a is F as being composed of concepts of *a* and *F* are the very same conditions it must meet in order to count as being about that state of affairs—so that a putative thought that a is F which failed to satisfy that condition could not count as being a *about* a's being F. But it is tempting to think that whatever we say about the connection between content and concepts for the case of thought will generalize to all kinds of states: if thoughts qualify as *representing* the world only by meeting a criterion that simultaneously qualifies them as representing the world conceptually, then any mental representation must meet this same criterion (hence, must satisfy the conditions required for being conceptual).

Yet if all of this is accepted, the project is in trouble. I shall show that the key argument that would support the restricted reading infers conceptual content from an

empirical datum of generality which is satisfied by all representational states if it is satisfied by any. Since generality in an organism's capacities to successfully discriminate distal features of the environment is the primary datum states with representational content are invoked to explain, all representational states exhibit Generality. Arguments for the existence of nonconceptual content based on failure to satisfy the Generality Constraint do not go through because they violate the very condition that would justify our regarding them as contentful mental states.

I believe this problem can be resolved by recognizing that representation in thought can be essentially conceptual without this being true for every kind of representational state. The essential connectedness of concept and content is anchored in facts about *what it is for a mental representation to be a thought*. The notion of a thought—of that kind of representational state—is anchored in capacities for inference, giving rise to *unrestricted* patterns of generality. Unrestricted generality is indirectly explained by the possession of concepts, inasmuch as it is directly explained by inferential capacities which rely on conceptual structure. Thoughts are essentially conceptual because *thinking*—the rational process which generate thoughts—makes essential use of concepts in the formation of states with objective content. Evidence for nonconceptual content would be the presence of a *restricted* form of generality in an organism's capacities to successfully discriminate distal features of the environment. It is possible—and even probable—that restricted generality is explained by different sorts of (non-inferential) processes through which subjects may acquire states with objective content. And if there are such processes, it is an open question whether they make essential use of concepts in the formation of states with objective

content. The possibility of nonconceptual content resides in the possibility of there being kinds of representational states which are not thoughts and which do not make use of concepts in their formation principles.

From this criticism of nonconceptualist proposals emerges a hopeful and more theoretically robust picture of what would justify attributions of nonconceptual content to mental states. This conclusion is anchored in a more theoretically satisfying picture of the explanatory value of ascribing conceptual contents to mental states. Although the conclusion of this chapter stops short of providing evidence for the existence of nonconceptual content, it provides clear guidelines for where decisive evidence for this conclusion will be found. Rather than looking for evidence directly in the surface features of representational states, we should use those surface features as a basis for exploring the kinds of processes which produce and utilize states with representational contents. Differences in the ways representational contents are formed ground differences in the kinds of contents that are so formed.

3.2 The Generality Constraint and Nonconceptual Content

I shall begin by quoting, at some length, the passage which leads up to Evans's formal statement of the Generality Constraint:

> (A)ny thought which we can interpret as having the content that a is F involves the exercise of an ability—knowledge of what it is for something to be F—which can be exercised in indefinitely many distinct thoughts, and would be exercised in, for

instance, the thought that b is F… And this of course implies the existence of a corresponding kind of ability, the ability to think of a particular object. For there must be a capacity which, when combined with a knowledge of what it is in general for an object to be F, yields the ability to entertain the thought that a is F, or at least a knowledge of what it is, or would be, for a to be F. And this capacity presumably suffices to yield a knowledge of what it is, or would be, for a to be G, when combined with a knowledge of what it is for an object to be G, for any arbitrary property of being G.

Thus, if a subject can be credited with the thought that a is F, then he must have the conceptual resources for entertaining the thought that a is G, for every property of being G of which he has a conception. This is the condition that I call 'The Generality Constraint.' (1982: 103-104)

The conclusion Evans reaches in this passage proposes a strong form of holism as being essential to thought: if a subject has the capacity to have a single thought (with a content of the form *that a is F*), they must have the capacity to have an indefinitely general range of semantically connected thoughts. More formally, the consequent condition refers to a thinker's capacity to think that range of thoughts that would be obtained by recombining the singular component of the thought ('a') with every 'property conception' the subject possesses (is G, is H, is J, etc). As is clear from the context, Evans intends for this constraint to be symmetrical: the thinker's capacities must equally extend to those thoughts that would be obtained by recombining the predicative component of the thought (is F) with

every 'particular conception' (a, b, c) the thinker has. The fact that Evans presents this constraint in the context of remarks about 'the nature of our conceivings' provides at least a partial explanation for it. Evans understands concepts to be cognitive abilities that are exercised in thought, and whose nature is spelled out through the consequent of the Generality Constraint. More specifically, a concept is a freely *recombinable constituent of a thought*. What it is for a thought to be 'conceptually structured' is for it to be "a complex of the exercise of several distinct conceptual abilities" (ibid., 100). Thus, for us to be justified in attributing a conceptual content to a thinker that thinker must possess the capacities described in the Generality Constraint—because possession of those capacities is constitutive (or partly constitutive) of a subject's possession of concepts. Amongst the many questions which Evans's discussion leaves unanswered, one is the question of why we should accept this characterization of conceptual content—and the constraint which is associated with it. Whilst this question will play a central role in the discussion to follow, it will be sufficient for our present purposes to note that many philosophers do accept it as valid (Peacocke, 1992; Dickie, 2001; Campbell, 1994; Tye, 2005; Carruthers, 2004, 2006; Camp, 2009; Beck, 2012).

The Generality Constraint can serve as a negative criterion for the existence of nonconceptual content if we can validly infer from the failure of perception to 'meet' or 'satisfy' or 'realize' the Generality Constraint that perception represents the world in a way that does not involve concepts. I shall focus on an argument for this claim as it appears in

Richard Heck's 2007 paper, 'Are There Different Kinds of Content?'[39] The first major premise in Heck's argument is a general assumption about what conditions would cause a kind of mental states to have a nonconceptual content. If we accept this as a valid criterion for conceptual content, we must also accept that representational mental state which does not satisfy it will have a content that is not conceptual. A crucial assumption in Heck's argument for this claim is the Generality Constraint is restricted specifically to a state's possession of conceptual content. It therefore allows for the in-principle possibility that there could be representational states which failed to satisfy generalizations of the same kind that cognitive states satisfy—and which would be states whose content is not conceptual. It is because of this, he stresses, that nonconceptual content cannot be ruled out on a priori grounds: if there is no such thing as nonconceptual content, as has been claimed by philosophers such as John McDowell and Charles Travis, this is a 'strong claim' that must be established on the basis of 'empirical' argument (2007: 11). In reading Evans in this way, Heck is in agreement with other philosophers who pursue the idea of nonconceptual content (Camp, 2009; Beck, 2012). The Generality Constraint allows, Beck insists, that mental states might fail to be systematic: "the Generality Constraint is a constraint on conceptual content, not a constraint on mental content *tout court*... For our purposes this is significant, since it means that if we can show that certain mental states violate systematicity we can use the Generality Constraint along with modus tollens to argue that those states have nonconceptual content" (2012: 565). If the Generality Constraint is a condition only on a state's possessing a conceptual content—and not on its

---

[39] Arguments for nonconceptual content along similar lines include Heck (2000), Peacocke (2001), and Tye (2005).

possessing a representational content as such—it is an open question whether a state could fail to meet that constraint whilst nonetheless counting as representational.

The second major premise in Heck's argument is that such conditions are in fact satisfied by perceptual states. Heck's discussion of this point centers on perceptual constancy; specifically, depth perception and color constancy. Heck claims that depth perception may be locally limited, so that relative distance from the subject is explicitly represented only for objects that are near one another in the visual field. Thus, for example, one object might be represented as closer than another that is but a short angular distance from it, and that object as farther away than another a short angular distance from it. But it could be the case that no such relationship between the first object and the third is explicitly represented at all. If this were the case, the Generality Constraint would not be satisfied, since it would not follow from the fact that one could perceive that a is closer than b and that b is closer than c that one could also perceive that a is closer than c. A similar phenomenon is illustrated in color constancy: one can visually perceive as explicitly uniform in color only small connected surfaces. If this were the case, claims Heck, the Generality Constraint would fail, "x is the same color as y" would be explicitly representable only under certain circumstances, for example, when x and y were points on a small connected surface. From the combination of this premise with the previous one, it follows that perception has a nonconceptual content.

One might seek to challenge Heck's conclusion by questioning the first premise. One might claim that even though systematicity is a valid criterion for a creature's possession of conceptual content in general, it need not follow that failures of generality constitute

evidence of nonconceptual content. Jerry Fodor is an example of someone who takes this kind of position: his view (which we shall consider more closely in the next section) is that even though generality is a feature of conceptual content in general, perceptual states have a conceptual content yet do not exhibit this generality because they are modularly encapsulated. Others have argued that it is not clear what the Generality Constraint is, and there is no acceptable construal of the Generality Constraint that makes non-trivially true both that conceptual states meet it and that perceptual states do not (Duhau, 2009: 41). Another reason to question this argument has its source in which Evans himself has to say about the Generality Constraint. The formulation Evans uses in his 'official' statement of the Generality Constraint articulates the antecedent condition as a subject's being *credited with a given thought*. Evans's discussion of the Generality Constraint certainly provides some support for the restricted reading of this constraint: in introducing the idea, he presents it as "a condition on the nature of our conceivings" (100), which might suggest that Evans is concerned uniquely with characterizing the conditions in virtue of which thoughts count as representing the world in a certain way—conceptually—whilst leaving untouched the broader question of what is involved in their representing the world full stop. Yet there are some parts of the discussion which suggest that Evans himself is not thinking of the Generality Constraint in the way nonconceptualists claim. Evans repeatedly references to the idea that the Generality Constraint holds for thoughts in a strong way: possession of conceptual capacities, as evinced by the satisfaction of the Generality Constraint, is a condition on thought *because it is a condition on thought-contents*.

It is a feature of the thought-content that John is happy that to grasp it requires distinguishable skills. In particular, it requires possession of the concept of happiness knowledge of what it is for a person to be happy; and that is something not tied to this or that particular person's happiness. There simply could not be a person who could entertain the thought that John is happy and the thought that Harry is friendly, but who could not entertain—who was conceptually debarred from entertaining— the thought that John is friendly or the thought that Harry is happy. Someone who thinks that John is happy must, we might say, have the idea of a happy man—a situation instantiated in the case of John (he thinks), but in no way tied to John for its instantiation. (ibid., 102. Italics are mine.)

We may illuminate Evans's remarks here by placing them in the context of Strawson's discussion of a similar constraint in *Individuals*. Strawson is concerned specifically with ascriptions of states of conscious experiences, for which he proposes an analog of the Generality Constraint: "it is a necessary condition of one's ascribing states of consciousness, experiences, to oneself, in the way one does, that one should also ascribe them, or be prepared to ascribe them, to others who are not oneself." He continues,

This means not less than it says. It means, for example, that the ascribing phrases are used in just the same sense when the subject is another as when the subject is oneself.

…

166

> The main point here is a purely logical one: the idea of a predicate is correlative with that of a range of distinguishable individuals of which the predicate can be significantly, though not necessarily truly, affirmed. (1969: 99)

The point Strawson is making here concerns objective reference; the conditions which must be met for a thought to be a thought *about that*.[40] The issue with states of consciousness in particular is that they are private. The property of being in pain is, by its nature, a property multiple people can have. That is part of what it is to be the property of being in pain. If your reference to pain is only in application to yourself, then in what sense are you referring to pain? Evans cites Strawson's remarks approvingly ("We should surely be reluctant to assign the content am in pain' to any internal state of a subject unless we were persuaded that the subject possessed an idea of what it is for someone—not necessarily himself—to be in pain.") In the Generality Constraint, Evans extends this argument to *all concepts*: "What we have from Strawson's observation, then, is that any thought which we can interpret as having the content that a is F involves the exercise of an ability—knowledge of what it is for something to be F—which can be exercised in indefinitely many distinct thoughts, and would be exercised in, for instance, the thought that b is F" (1982: 103). What he is saying, rather, that a subject's being related to representational contents which are semantically general in these ways entails a further, fundamentally *psychological* generality—which is part of what it is for a *subject* to possess a

---

[40] This theme goes all the way back to Frege, who as we saw in chapter one rejects 'ideas' as ways of capturing our thoughts because they are not sharable: "Is that lime-tree my idea? By using the expression "that lime-tree" in this question I have really already anticipated the answer, for with this expression I want to refer to what I see and to what other people can also look at and touch" (1918/1956: 300).

concept. The idea is that if a subject has a thought with a conceptual content—one which refers to objects and properties by means of that kind of semantic structure, they must grasp *what it is for a thing to be F*. Entertaining a thought with a predicative structure entails that one's understanding of which property is attributed to the object is entirely independent of one's knowledge of what it is for any particular thought to be true: it must be independent of one's knowledge of what it is for the property to be instantiated in any particular case. And this knowledge is manifested in their capacity to recombine the constituents of their thoughts: to grasp what it would be for a thing to be F is to have the capacity to represent *b* as F (or c, or d) just as well as one can represent *a* as F.

All of these worries will be of central concern in the remainder of this chapter. My aim in the next section is to draw out a picture of the picture of concepts that lies behind the restricted way of reading the Generality Constraint. In sections 3.4 and 3.5 I shall argue that this picture cannot serve the project of establishing nonconceptual content in the ways its defenders take it to; and that it is moreover a picture we have independent reasons to reject. In section 3.6, I shall turn to the task of providing an empirically adequate argument for the existence of conceptual content; and I shall show how this argument justifies a global reading of the Generality Constraint. In section 3.7, I'll show that far from ruling out the possibility of nonconceptual content, this argument may actually support it.

3.3 The Argument from Systematicity

If nonconceptual content is indeed a theoretical possibility this must be shown to flow from a correct understanding of what's involved in the conceptuality of thought. For the

philosophers we are considering, the justification for the Generality Constraint lies practices of empirical psychological explanation. Constraints on ascriptions of representational contents reflect the kinds of properties thought-contents are invoked to explain. Likewise, the Generality Constraint—a constraint on ascriptions of conceptual contents—reflect facts about what that *kind* of content is introduced to explain. The Generality Constraint is a constraint *only on* conceptual content and not on mental content tout court because concepts are introduced to explain facts that are logically distinct from those which contents are introduced to explain.

Representational contents are not entities that we have independent access to: one could not hope learn whether a person has beliefs by *looking for them*—for, even if one were equipped with an FMRI of their brain, one would not find it. Why should we attribute content to mental states at all? Consider Heck's answer to this question:

What would we lose if we just ignored [the representational properties of mental states]? I take it that we would lose the very idea of psychological explanation. We are in the habit of explaining our own behavior, and that of other creatures, in terms of what we all believe: We explain why Bob ran across the room in terms of his believing that his stuffed dinosaur was on the other side. These explanations are typically causal and counterfactual supporting, which is to say that there is a law of one sort or another that, if a given explanation is correct, it instantiates. The explanations themselves are formulated not in terms of the neurological features of mental states but in terms of their contents, and the same is true of the laws. And so

we might say: The reason we should attribute content to mental states is because there are things we wish to explain in terms of mental states, as individuated by their contents. (2007: 6)

We believe there are such things as representational contents even though we cannot see, smell or otherwise sense them, because they contribute to practices of *genuine* explanation. Consider the following claim: 'Croesus believed that the gods had determined that he would win the war against the Persians.' We have reason to think that claims such as this one are true to the extent that they give us a unified explanation for various facts. It seems to fit with the fact that Croesus consulted the Oracle of Delphi who communicated an ambiguous message that could have been interpreted in this way. And it allows us to understand why Croesus, who had nothing to gain and everything to lose from a war with the Persians, embarked on such a war, with disastrous consequences for himself and his nation. Without this level of explanation— explanations which invoke a subject's representational perspective on the world—we would lose our ability to make sense of animate action. As Simon Baron-Cohen has pointed out, the 'mindblind'—one who is "aware of physical things, but blind to the existence of things like beliefs, knowledge, desires, memories, and intentions"—has an impoverished capacity to explain actions as compared with a "mind-reader" (1996: 2). Even a very basic sequence of acts—a person's walking into the bedroom, walking around, and then walking out again—is a real mystery to such a person. It is primarily because of their role in explanations of this kind that we have reason to think some form of realism about propositional attitudes is true: we have

reason to think there are such states because without them, we would lose a distinctive level of explanation for why subjects act as they do.

Just as we don't have direct perceptual access to the existence of thoughts, so we don't have direct perceptual access to the existence of concepts. Why, then, should anyone think there are such things as concepts? Following the foregoing observations, we may assume that the answer lies in their role in genuine psychological explanations. Yet many have thought that this explanatory role must be something over and above the considerations which motivate ascriptions of content per se. For we can imagine a creature that represents that John is happy and that Harry is sad, but whose ability to think each one is utterly independent of their ability to think the other. As Camp observes, it is tempting to feel that these thoughts could still perform the basic job of belief— that of explaining the creature's behavior—"so long as they were formed and extinguished in the appropriate circumstances, and produced appropriate further actions when co-instantiated with other attitudes." In that case, though, there would "be no point in attributing concepts to that thinker; we could simply describe its thinking in terms of whole beliefs about entire states of affairs" (2009: 277). Beck agrees: "that Amy is funny and that Bob is gentle, but that, for whatever reason, completely lacks the capacity to represent that Amy is gentle" (2012: 565). The reason we regard thoughts as conceptually structured must be because they perform an additional "job" or have a "point" *beyond* their basic function of explaining action.

One widely-endorsed proposal for what that role might be is found in 'Connectionism and Cognitive Architecture,' in which Fodor and Pylyshyn articulate what

has since become the classic cognitivist defense of concepts. The argument from systematicity is articulated as follows:

> What does it mean to say that thought is systematic? Well, just as you don't find people who can understand the sentence 'John loves the girl' but not the sentence 'the girl loves John,' so too you don't find people who can think the thought that John loves the girl but can't think the thought that the girl loves John.… But now, it the ability to think that John loves the girl is intrinsically connected to the ability to think that the girl loves John, that fact will somehow have to be explained. For [a representationalist] the explanation is obvious: Entertaining thoughts requires being in representational states. And… the two mental representations… must be made of the same parts. But if this is right, then mental representations have an internal structure. (1988: 39)

The argument has two major premises. The first is that thought is as a matter of fact systematic, or equivalently that thought 'realizes' or 'satisfies' the Generality Constraint: a thinker who is capable of entertaining the thought that a is F and is also capable of entertaining the thought that b is G will typically also be capable of entertaining the thought that a is G and that b is F. The second premise is that *if* thought is systematic, then it must be composed of concepts—the idea being that unless the thought that a is F were composed of recombinable constituents, a concepts of a and a concept of F, the systematicity of thought would be an inexplicable mystery. It is possible to read Evans's

Generality Constraint as formalizing this argument for the existence of thoughts. The presence of generality in a subject's thought is our starting data: we 'observe' that it is realized by thoughts; The claim that the contents of beliefs are structured is motivated by the fact that they 'satisfy' the Generality Constraint as as an inference to the best explanation for what we observe. Heck reads Evans in this way: "We want not only to observe that there is a certain pattern in people's ability to entertain various thoughts, we also want to explain this fact. The explanation Evans suggests is that the capacity to entertain the thought that a is F has a structure that corresponds to the structure of the thought itself: … The ability to think such a thought thus depends upon, and is made possible by, one's ability to think of a and to think of an arbitrary thing that it is F" (Heck, 2007: 9).

These parallel arguments for the existence of representational contents and concepts form the basis for a pair of *criteria* governing our ascriptions of contents and concepts. We are justified in ascribing the content *that a is F* if this content makes a genuine contribution to explaining a creature's behavior. Ascriptions of personal-level contents are answerable to the overarching constraint of making the subject of the attributions intelligible (Peacocke, 1995: 238). The validity of the constraint for determining content in a given case derives from the general conditions that govern ascriptions of representational content in general. We are justified in holding that the content that a is F is conceptually structured if and only if this subject's thoughts (or representations) are systematic in ways that would be usefully explained by conceptual abilities. It is because concepts *contribute* to explaining the satisfaction of the Generality Constraint that a subject

must satisfy the Generality Constraint to count as having concepts. The crucial point to note, in the present context, is that these criteria are conceptually distinct. The generality (or 'systematicity') exhibited by conceptual thought has nothing to do with their basic function of explaining a subject's behavior. That is why possession of a conceptual content is merely a contingent or incidental aspect of thinking about the world: something a representational state could lack without ceasing to be a representational state. There is space for the possibility of nonconceptual content because there exists an explanatory gap between a mental state's representing the world and its representing the world conceptually. We have traced the explanatory contributions of contents and concepts to two observable properties of thoughts which, although they happen to be co-instantiated in the case of thoughts, are logically distinct. Whether these two properties are co-instantiated in every case is not the sort of question that could be settled from 'the comfort of the philosopher's armchair.' In particular, there is no conceptual basis for ruling out the possibility that there could be a creature whose behavior could be explained in terms of intentional states yet which did not possess the kind of systematicity in their thoughts that would license attribution of conceptual content. And if we were to encounter such a creature, the constraints imposed by practices of psychological explanation would seem to dictate that this creature possessed mental states with a representational content which was not conceptually structured—since ascribing concepts to such a creature would play no role in explaining facts about such a creature. It is an empirical question whether there could be one of these properties in the absence of the other; and as such, it is also an

empirical question whether there could be a creature who possessed states with a representational content which was not composed of concepts.

3.4 The Problem

The datum of systematicity evokes wildly different responses amongst philosophers. Those who treat systematicity as the primary basis for our ascriptions of concepts—holding that concepts are directly read off the surface structure of thoughts—take systematicity to be self-evidently true. For example, Camp claims that "(e)ven if one denies that the Generality Constraint follows ineluctably from the very nature of thought, something like the requisite generality clearly applies to our thinking, and differentiates it from the mental representings of other animals" (2004: 210). And according to Fodor and Pylyshyn the systematicity of linguistic capacities is a "*fact* no-one can plausibly deny." Yet apparently not everyone is so optimistic. Kent Johnson has contested the datum outright: "I do not think it is very clear what systematicity is, and I am even more skeptical that language (or thought) has it" (2004: 112). And Evans himself comments that generality is an 'ideal' to which our actual system of thought only approximately conforms (1982: 105).[41]

---

[41] The only practical example Evans provides in his discussion of the Generality Constraint further counts against the idea that he takes Generality to be an empirical *datum* which concepts explain—rather than an idealization. According to Evans, most of us would be "reluctant to assign the content am in pain' to any internal state of a subject unless we were persuaded that the subject possessed an idea of what it is for someone—not necessarily himself—to be in pain, and unless we were persuaded that the internal state in question involved the exercise of this

Idea" (1982: 103). Evans then goes on to claim that the Generality Constraint is a consequence of this intuition. Assume, for the sake of argument, that Evans is right in claiming that this is an intuitive criterion we hold thinkers to as a matter of practical fact. A crucial thing to notice is that this is a far distance removed from the recombinability condition. The criterion illustrated by this example is negative: what we need to know, in order to feel justified in our ascribing the content 'I am in pain,' is that the ability should not be tied to one particular thing. The recombinability condition is positive: it specifies a (wide) range of thoughts the

The explanation for these differing points of view lies in the fact that systematicity is not a well-defined phenomenon. To be sure, we have the intuition that there is some kind of generality that is inherent in the idea of a concept: whereas a subject who lacks a concept of Christianity can't think any thoughts about things which are Christian, having acquired that concept they possess a kind of knowledge that can be used in thinking about a range of different kinds of things. However, systematicity says something stronger than this: that a possessor of concepts has a capacity to recombine their concept with *any* arbitrary singular concept. And this is much less obvious. There are so many caveats that one might reasonably question how valuable the remaining thesis is (to quote Monty Python's Life of Brian: "All right, but apart from the sanitation, the medicine, education, wine, public order, irrigation, roads, a fresh water system, and public health, what have the Romans ever done for us?") Robert Cummins has criticized the systematicity intuition on the grounds that it illicitly trades on facts about the sentential tokens in terms of which the claim is expressed: it seems undeniable that anyone who can think that a is F and b is G is also able to think that a is G and b is F because the latter formulae are straightforward permutations of the former; but the intuition disappears if we replace those expressions with synonymous sentences which do not involve permutations of the expressions involved.[42] The datum is something that needs to be clearly characterized and established through argument.

---

subject must be capable of thinking. You could meet the negative condition without meeting the strong positive condition.

[42] Cummins's argument for why the systematicity of thought depends on linguistic expression is as follows. Consider these two statements: (1) Anyone who can think that John loves Mary can think that Mary loves John; (2) Anyone who can think Mary's favorite proposition can think that Mary loves John. (1) is intuitive because of the fact that 'John loves Mary' is a permutation of 'Mary loves John.' But the intuitive force of (1)

Establishing systematicity through language: derive the systematicity of thought from the systematicity of language. The first premise in the argument is that language is systematic. Fodor characterizes linguistic systematicity as follows:

> The property of linguistic capacities that I have in mind is one that inheres in the ability to understand and produce sentences. That ability is—as I shall say—*systematic*: by which I mean that the ability to produce/understand some of the sentences is *intrinsically* connected to the ability to produce/understand many of the others. (1987: 149)

The other major premise is that if language is systematic then so is thought; so if one accepts the systematicity of language, one must also accept a corresponding systematicity at the level of thought. Thus Fodor and Pylyshyn write,

> (I)n the case of verbal organisms the systematicity of thought follows from the systematicity of language if you assume—as most psychologists do—that understanding a sentence involves entertaining the thought that it expresses; on that assumption, nobody *could* understand both the sentences about John and the girl unless he were able to think both the thoughts about John and the girl. (1988: 39)

---

disappears if we replace it with (2). On the obvious assumption, the italicized phrase in (2) refers to the same proposition as the italicized phrase in (1); hence (1) and (2) express the same claim. But, as Cummins points out (2) elicits no systematicity intuitions precisely because the italicized phrase in (2) is not a permutation of 'Mary loves John.' "It is the structure of the mediating representation which determines whether or not we see systematicity in the thoughts" (1996: 596).

An initial reason to worry about basing systematicity in language, particularly if we are interested in nonconceptual content, is that any non-linguistic states will turn out on this view to be *eo ipso* nonconceptual. But this seems too easy. Since linguistic capacities imposes requirements going above and beyond the possession of conceptual capacities, it seems in principle possible for a creature's mental states to be conceptual even if they are not linguistically expressible (perception would be a case in point). In addition to this worry, the argument is problematic considered purely on its own terms. Fodor and Pylyshyn claim that the systematicity of linguistic capacities is a "*fact* no-one can plausibly deny." Yet many have denied it. Johnson argues that linguistic systematicity is a generalization with so many exceptions that what remains is less than impressive, if useful at all. Moreover, even if language were systematic, there are problems with the assumption that we can straightforwardly carry conclusions about language over to the case of thought. Linguistic abilities could be systematic, even if representational contents were not (Stalnacker, 1998). Based on these worries, I think we should reject attempts to ground the systematicity of thought in language.

The other main proposal for how to justify the systematicity of a creature's cognitive capacities is to infer these from a more basic kind of systematicity in a creature's capacities to discriminate and respond to objects and properties in their environments. Fodor assumes this connection when he claims that if cognition were not systematic,

> [i]t would have to be quite usual to find, for example, animals capable of learning to respond selectively to a situation such that aRb , but quite unable to learn to respond selectively to a situation such that bRa (so that you could teach the beast to choose the picture with the square larger than the triangle, but you couldn't for the life of you teach it to choose the picture with the triangle larger than the square). (Fodor, 1987: 153)

Unpacking this proposal, the initial datum is not systematicity of cognitive abilities but a more basic kind of systematicity in a creatures discriminatory abilities. For example, if a creature can discriminate and appropriately respond to the presence of a green ball, and likewise for a yellow banana, they are typically also capable of discriminating and responding appropriately to a yellow ball and a green banana. The best explanation for the fact that a creature is capable of responding to systematically related *states of affairs* is that the creature has the capacity for systematically related thoughts (e.g. beliefs). Once this conclusion is established, the idea is that one can then apply the standard argument from Systematicity to establish the conclusion that the creature represents the world conceptually. What justifies the move from behavioral capacities to representational states? The argument from 'aspect' (which Michael Tye also endorses): "Our having seen a cat moving in the bushes may well be part of the story about why we turned towards the bushes; *but not unless we saw the cat as something moving in the bushes.* What enters into the explanation of our behavior isn't what we see saw per se; it's what we see what we saw *as.* Take concepts away, and psychology loses seeing as. Take seeing as away, and it loses the junction between perception and action (2015: 211). The theoretical problem

sometimes called 'the problem of underdetermination,' is that the registrations of causal effects on sensory organs underdetermines the distal features of the environment that cause them. Subjectivity of the senses: register only how the world *affects us.* The same fact can affect us, or 'show up for us,' in myriad different ways. And how things affect us can mean different things at different times. Each of the different possibilities may have quite implications for a creature: some are to be avoided, others to be taken advantage of. Some may be a matter of life-or-death. Where there is underdetermination, we can't appeal to the stimulus inputs to explain the response; hence, we must appeal to an internal representation to account for why the subject responds in one way, rather than in some other way. Which means, therefore, that there is an explanatory work for conceptually structured thoughts to do.

The argument from discriminatory responses is grounded in a datum that is uncontestable; and it avoids the problematic consequence of the argument from linguistic systematicity in providing a barometer for conceptual capacities that is independent of linguistic abilities. Yet where linguistic systematicity is too strong a criterion, behavioral capacities turn out to be too weak. Whilst it is tempting to feel that if there is successful engagement with the world there must be an 'interpretation' of the stimulus which explains this, as a general rule this is incorrect. For compare the following scenarios:

1. Over breakfast with my partner, I utter the sentence, "What's the weather today?" In response, my partner utters the sentence, "It's going to be really hot today. I think they said 87 degrees."

2. Over breakfast with my partner, I utter the sentence, "Alexa, what's the weather today?" (Alexa is our electronic 'home help,' having the capability to issue whether reports, make shopping lists, set timers, and so forth, all via voice commands). In response, Alexa utters the sentence, "Today, in Darien, you can expect sunny weather, with a high of 87 degrees, and a low of 70 degrees."

From the point of view of their observable features, the two situations seem to be more or less the same: the only difference is who (or what) is doing the uttering. Yet, intuitively, they have different underlying explanations. Here is a natural explanation for what is going on in the first scenario: my partner interprets my utterance as having the meaning [*what is the weather forecast today, here where we are // form of a interrogative]*; he believes on that basis that I wish to know what the weather will be like today; and this belief, combined with a desire to oblige my wish, causes him to respond as he did. Here is a natural explanation for what is going on in the second scenario: Alexa is fitted with sensors which detect patterns of sound and a speaker which generates sounds: she has been programmed, given certain patterns of sounds as inputs, to respond to certain patterns of sounds (corresponding to speech phonemes) the same form of explanation is not a natural explanation for the second scenario. It is true that sometimes, my partner and I will use psychological terms to describe Alexa's behavior, but we mean this: we do not need such notions as *interpretation*, *belief*, or desire to explain Alexa's functioning. Alexa has a function that capitalizes on this dependence. Alexa illustrates a more general phenomenon which Tyler Burge has called functional information-carrying (2010). A mechanism is

designed, or evolves, which brings about an automatic response in a system to the proximal state; where this response serves a function that is tied to the distal state. An information carrying state will trigger a response, in the system, that serves a function. The important thing to notice is that, although it is possible to describe Alexa, and other systems of this kind, can be described *by mean of thought*, thoughts are not part of the nature of the process. They are merely a gloss or afterthought on a process which is fully explained in non-psychological terms. Given that the same mechanisms which underlie Alexa's functioning are widespread amongst the capacities of sentient organisms, the need to distinguish genuine representation from information registration is no mere idle worry. There are many examples of functional information carrying systems which have been 'programmed' by evolution or acquired through associative learning. The link to the environment is secured in a purely 'external way,' through functions that 'latch on to' the laws of nature. Their explanation includes reference to causation, and to biological function; but they does not include any psychological entities. Hence, creatures may have systematic capacities to systematically respond to features of the environment which is not underwritten by thought. If systematic behavioral competencies does not succeed in distinguishing states which have the function of representing the world from systematic abilities which are underwritten by 'mindless' cases of information registration, it does not succeed in showing that these are cases of genuine (rather than pseudo) explanation. Trivially, then, it does not succeed in establishing the presence of concepts.

3. Establishing systematicity through stimulus independent behavioral competencies. Consider the following passage, which describes the well-documented pattern of behavior in ants to dispose of the corpses of their fellow ants:

> When a corpse of a Pogonomyrmex barbatus worker is allowed to decompose in the open air for a day or more and is then placed in the nest or outside near the nest entrance, the first sister worker to encounter it ordinarily investigates it briefly by repeated antennal contact, then picks it up and carries it directly away towards the refuse piles. . . . It was soon established that bits of paper treated with acetone extracts of Pogonomyrmex corpses were treated just like intact corpses . . . the worker ants appear to recognize corpses on the basis of a limited array of chemical breakdown products. They are, moreover, very narrow-minded on the subject. Almost any object possessing an otherwise inoffensive odor is treated as a corpse when daubed with oleic acid. This classification extends to living nestmates. (E. O. Wilson, 1971; quoted in Allen and Hauser, 1991)

Witnessed under normal circumstances, the behavioral responses of ants might easily be mistaken for cases of thought; one might even be inclined to feel that one must have underestimated the intellect of ants; that they must somehow be 'smarter' than one has suspected. Yet in the light of their inability to vary their responses to oleic acid even in the presence of obvious evidence to the contrary, we have the strong intuition that genuine

thought is somehow importantly different from the capacities of ants. Describing these differences, Allen and Hauser write that,

Like ants, humans may be duped into believing that someone is dead when they are not (e.g., by the effects of some drugs). But unlike ants, humans who have been duped in this way normally will modify what they take in the future to be evidence of death so as to be careful not to be fooled by appearances. Humans are capable of recognizing something as dead because they have an internal representation of death that is distinct from the perceptual information that is used as evidence for death. … We would attribute an abstract concept to an organism if there is evidence supporting the presence of a mental representation that is independent of solely perceptual information. (1991: 231)

The disposal of corpses confers a selection benefit for both ants and humans, protecting communities from the spread of disease. Yet the description of the conditions under which the ants *fail* to accurately discern dead ants make it obvious, somehow, that they are not thinking. The kinds of mistakes they make feel somehow 'odd.' At the center of the phenomenon is a causal relation between one kind of property, to which ants are directly sensitive (Oleic acid), and another kind (dead ant) to which they are not. There is a 'law of nature,' to the effect that when ants decompose, they emit the chemical, oleic acid. The problem arises because the reverse does not hold: the presence of oleic acid is normally a highly *reliable* indicator of the presence of dead ants; yet there may be oleic acid present

without the presence of a dead ant (for example, if a scientist drops this onto a piece of paper). The key feature of this example which accounts for our intuitions that what we have is not a case of genuine representation is that the ants are insensitive to contextual cues which clearly demonstrate that this is one of those aberrant cases. Faced with perceptual evidence that a putative corpse is moving (or is a piece of paper) a thinker would not judge (or would revise their judgement) that what they are confronted with is a corpse. Yet the ants are completely insensitive to any such evidence. One might hold that our intuitions in these cases serve as a useful guide to understanding what sorts of considerations are relevant to distinguishing genuinely conceptual thought from mindless or automatic processes of information registration. Genuine thought requires 'stimulus-independendence.' It requires some distance between the representation and what it is a representation *of*; "(o)therwise, the world, and not the thinker, is shouldering the bulk of the representational burden. And if this is so, then that ''thinker'' really is just a passive reactor" (Camp, 2009: 288). Rather than counting any kind of systematicity as evidence of conceptuality, a practical fulfilment of the Generality Constraint should reflect the fact that truly conceptual thought is relatively independent of a thinker's current circumstances.

The problem here is that, if we accept this way of characterizing systematicity, there is no hope of using the Generality Constraint as a criterion for the existence of states with nonconceptual content. Evidence of discriminatory capacities which are both systematic and stimulus-independent is a minimal condition for us to be justified in counting a state *representational* in any sense. What Heck fails to acknowledge, in his argument from perceptual constancy, is that constancies constitute a form of stimulus-independence in

perception which qualifies perception as a representational state. The range of stimulus-conditions which may give rise to a perception of shape is, as Ruth Millikan points out, quite impressive:

> Think of the variety of proximal visual stimulations—what actually hits the eye - to which a given shape may give rise when viewed from various angles, from different distances, under different lighting conditions, through various media such as mist or water, when colored different ways, when partially occluded… And shape is coidentified by the haptic systems. (Eye-hand coordination is a nice example of coidentification through concurrent rather than serial sources of information.) You can feel the shape of a small object in your hand in a variety of ways, for example, with these fingers or those, when the object is turned this way or that way, perhaps by using two hands, by merely holding the object or by actively feeling or stroking it. You can perceive larger shapes (say, in the dark) by exploring with larger motions that involve your arms, body and perhaps legs, and by employing the touching surfaces of a wide variety of your body parts. (2012)

The argument for nonconceptual starts from the premise that perception 'violates' the Generality Constraint: our perceptual abilities lack the kind of systematicity that is supposed to be constitutive of conceptual representation. Yet once systematicity is made precise, perception is much more like thought than we might have supposed. According to the argument of the previous section, the datum of systematicity consists in evidence of

discriminatory capacities which are both systematic and stimulus-independent. Any conceptual state *must* exhibit such features if it is to count as conceptual, because these are the features of thought which motivate us to introduce the notion of a concept in the first place. So a violation of the Generality Constraint will consist in a failure of representational states to correspond to stimulus-independent representational competencies. Yet stimulus-independence is a feature of all and any representational states—including perception. Generality is not, as these philosophers would have it, logically distinct from the data which motivate representational content: if stimulus-independent behavioral competencies are a sufficient basis for directly inferring a creature's possession of concepts, then any state which counts as representational will automatically qualify as conceptual.

3.5 Degrees of Stimulus-Independence

A possible response to the line of argument I've been developing against the nonconceptualist is suggested by Elizabeth Camp. Camp's 2009 paper, 'Putting Thoughts to Work,' argues that the cognitive states of some nonhuman animals count as having a nonconceptual content because they fail to satisfy the Generality Constraint in a 'robust' way. In arguing for this conclusion, Camp deploys the criterion of generality in a way that shows sensitivity to the worry I've been pressing. She recognizes the need to distinguish genuine representation of the environment from purely mechanical 'passive reaction': a practical fulfilment of the Generality Constraint should reflect the fact that truly conceptual thought is relatively independent of a thinker's current circumstances. However, her agreement with the intellectualist is tempered by a recognition that stimulus-independence

comes in degrees. Limited forms of generality in creatures with basic cognition that, whilst not *conceptual*, are nevertheless *representational*. For example,

> Honeybees can represent states of affairs that they have never actually encountered, and that are distant from them in a quite literal sense of the term. A wide range of animals can represent properties at distant locations, and navigate to those locations by novel routes to satisfy their desires... Rats can navigate to remembered locations without relying on any local landmarks, for instance by swimming to a submerged platform in an opaque pool of water (2009: 290).

Camp concludes from these observations that we do have good reason to describe creatures with just basic cognition as capable of transcending their current circumstances to represent absent situations, and not merely as responding differentially to immediate stimuli (ibid., 289). However, the restrictions on which thoughts creatures with mere basic cognition can think in a given situation demonstrates that although they might satisfy the Constraint "in principle" (given the appropriate stimulus inputs), they do not possess the breadth of stimulus-independence that is a 'practical condition on satisfying the Generality Constraint in a robust way. The generality their thoughts possess, such that it is, is nevertheless restricted in that it is dependent on the world presenting the creature with the appropriate stimuli. For this reason, these creatures have nonconceptual thought, or 'basic cognition,' as distinct from the "active, self-generated cognitive flexibility" possessed by creatures with capacities for conceptual thought.

The idea that systematicity comes in degrees seems to provide an attractive response to the problem I have described for nonconceptualist arguments regarding perception. There are those who have thought that perception is comparable, in its generality, to capacities for thought. Ruth Millikan, for example, suggests this point in the context of an analogy between conceptual recognitional abilities and perceptual constancies (which she groups together under the heading of 'unicepts'). She compares the generality of perception with with a vivid list of some of the ways you might be able to recognize a person, Bert:

> You can do this by seeing Bert in the flesh, 20 meters up the street, ... by recognizing Bert's signature or handwriting, by recognizing Bert's style of prose or humor or, perhaps, of musical interpretation or of some other activity, by the sound of the instrument Bert plays coming from the next room or the hammering that accompanies Bert's current home project, by recognizing Bert's name when someone speaks it, or when it is written, by hand or in any of a hundred fonts, and so forth. Also, surely, you can recognize that the information arriving is about Bert through many hundreds of descriptions of Bert, the person who was or did this or that, about whom this or that is true, or you may recognize whom the information is about using various kinds of inference, induction or abduction (2012).

Both of these phenomena she takes to be underwritten by 'unicepts' (a variant of a concept). Yet whilst the presence of stimulus-independence marks the discovery of

widespread and surprising similarities between perception and thought, these discoveries do not negate the differences we started out by describing. Perception is responsive only to certain sorts of stimuli, under certain sorts of conditions. You can recognize Bert in indefinitely novel and open ended ways, but you can't see colors in the dark or perceive the shape of something whose edges are obscured from sight. It seems bad practice to overstate the degree of the similarity between perception and thought at the expense of these significant differences,

Ultimately, I'm going to argue that there is something importantly right in this idea that different degrees of systematicity are relevant to grounding distinctions between conceptual and nonconceptual kinds of content. But the point I want to stress at present is that this idea isn't something we can capitalize on within the framework provided by the argument from systematicity. The problem can be traced back to the assumption that concepts are somehow supposed to be directly inferred from the surface structures of thought. The idea is that the atomic structure of concepts somehow directly explains the patterns of generality that are visible on the surface of thought. Thinkers entertain thoughts within the ranges they do *because* their thoughts are composed of constituents whose structure permits subjects to freely recombine them. The problem we have been considering is that giving a non-question-begging description of the kinds of surface structures that would licence this inference is remarkably hard to do. There is no perfect systematicity in thought, as there might be in a simplified computer language. Thought seems to be messier than a logical analysis of its constituents would suggest; shot through with its own kinds of context-dependence which aren't neatly captured by the model of a

'word-like' constituent. This messiness is accounted for by introducing constraints from outside of content: But once we allow for the idea of some degree of imperfection in the systematicity of thought, it is hard to see how we can draw a principled distinction between that kind of case and a case which involves "too many" limitations. It just seems arbitrary to insist that the limitations we observe in the case of perception are due to the nature of perceptual content whilst the limitations we observe in the case of thought. This is where the nonconceptualist plays right into the hands of their conceptualist rivals. Thought is systematic, except for the cases which are due to 'psychological factors' or 'modularity' or 'fragmentation.' Why can't we say the same thing for perception—the only difference being that these externally-imposed constraints are greater for perception than for thought (perhaps because of perception's function as an input system)?

What conclusions shall we draw from this? The conceptualist takes it as evidence that all representation is, after all, conceptual. But I'm inclined towards another possibility: that the very framework within which we've been considering these questions is at fault; that the problem lies in trying to derive the idea of a concept directly from the systematicity of thought. Notwithstanding the frequent claims that the argument from systematicity constitutes an empirically-oriented approach to understanding the nature of concepts, in my view there's something deeply unsatisfying about this argument— something *empirically* unsatisfying. In the worst case, it presents the conditions under which ascriptions of concepts are justified as ones that are unrecognizable from our commonsense observations about how thoughts behave. As Mark Wilson has said, "We will have told the story of concepts wrongly if it doesn't turn out to be one where our usual

forms of conceptual evaluation emerge as appropriate and well founded most of the time" (2008: 3). In the best case scenario, it leaves the empirical foundations of concepts fundamentally vague. In the next section, I want to suggest a different way of thinking about the empirical basis for ascribing a conceptual content to thought. In the following section, I'll show that this also gives us a satisfying basis for understanding the explanatory role of nonconceptual content.

3.6 The Argument from Inference

There are two parts to the argument. The first part establishes a link between the stimulus-independence of thoughts and their representational nature. The central claim is that,

> The patterns of stimulus-independence we observe in thought provide support for the existence of states with representational contents insofar as they are explained by inferential processes which utilize and produce states with representational contents.

The second part establishes a *parallel* link between the stimulus-independence of thoughts and their conceptual nature. The central claim is that,

> The patterns of stimulus-independence we observe in thought provide support for the existence of states with conceptual contents insofar as they are explained by

inferential processes which utilize and produce states with conceptually structured

representational contents (contents composed of freely recombinable constituents).

I shall spell out the support for each of these claims in detail momentarily. But what I want

to point out initially is that this general argumentative strategy for the conceptuality of

thought weaves together the same sorts of phenomena we have been considering all along

in a fundamentally different sort of explanatory pattern. The high degree of simulus-

independence we find in thought provides evidence of the conceptuality of thought

indirectly, by providing direct evidence of inferential capacities which essentially depend

on concepts. The Generality Constraint is a valid theoretical criterion for a thinker's

possession of concepts not as a mere reflection of a corresponding generality that we

'observe' in a creature's thoughts, but because it articulates a general principle about what

concepts *must be like* both for them to make possible inferential capacities and for them to

be the products of such capacities.

The first claim establishes a link between the contentful nature of thought and a

thinker's inferential capacities. To motivate this link, it will be helpful to go back to the

intuition we have about stimulus-independence. Earlier I used the example of the

responses of ants to oleic acid to illustrate this intuition. Another example, discussed by

Temple Grandin in her book, 'Thinking in Pictures,' concerns a severely autistic man, Ted

Hard who

has almost no ability to generalize and no flexibility in his behavior. His father describes how, on one occasion, he put wet clothes in the dresser after the dryer broke. He just went on to the next step in a clothes washing sequence he had learned by rote. (2006: 34)

If we were to witness Ted's behavior when things are operating normally, we might be inclined to ascribe thoughts to him. But the context reveals Ted's responses to be too dependent on a given stimulus—his perception of the end of the dryer's cycle—to be explained by a *belief* that the clothes are dry. Why should that matter to our assessment of the nature of his underlying mental states? For one thing, the fact that his response correlates so closely with the perceived end of the dryer cycle draws our attention to a possible non-intentional explanation for his behavior: clothes being dry is correlated reliably *enough* with the end of the dryer cycle that one might reasonably develop an automatic or 'instinctive' association between the two. But this can't be the whole story. For we allow that thinkers may make mistakes—miss key evidence—in reaching conclusions about their environments without thereby revoking our ascriptions of contents to those subjects. The problem in this case seems to lie in the nature of Ted's mistake. The evidence is *too obvious* for him to have missed it. If Ted can't take the sopping wet sensation of the clothes in his hands as a reason *against* the proposition that the clothes are dry, he can't have been taking the end of the dryer cycle as a reason *for* that proposition. This suggests that our intuitive criteria for ascribing thought to a subject are somehow tied to that subject's possession of inferential capacities.

The stimulus-dependence of certain responses to environmental entities counts against their being representational by revealing the nature of the process which explains the creature's successful engagement with the environment. The organism has a capacity to track the feature in question. The patterns of success and failure inherent in that capacity are best explained by a certain sort of causal process occurring within that organism: functional information-registration. Such a process carries determinate predictions concerning when the creature will succeed (when the stimulus is present) and when the creature will fail (when the stimulus is not present)—and it counts as a sound explanation inasmuch as its predictions match what we observe. By being connected to a particular kind of process in this way, stimulus-dependence licenses the conclusion that the underlying states of the organism are not representations of the feature in question. For the process is purely mechanical and non-psychological. At no point do we require the idea of a state with representational content to account for that process. Such an attribution would be explanatorily superfluous. Stimulus-dependent states are not thoughts (or representations of any kind) because the idea of a representational content has no role to play in the explanation of their capacity.

In an entirely parallel way, then, we can say that the stimulus-independence of thought counts in favor of their being representational by revealing the nature of the process which explains the creature's successful engagement with the environment. Again, the organism has a capacity to track some feature of the environment. Yet the patterns of success and failure in the organism's responses to that feature reveal that a fundamentally process underwrites their capacity: inferential processes. Such a process predicts when the

creature succeeds in terms of factors such as their recognition of evidence which supports a conclusion and their possession of background knowledge which licences a connection between the evidence and the conclusion. And it predicts when the creature will fail. Attributions of inferential forms of problem-solving are valid inasmuch as the patterns of success and failure predicted by this process are borne out in the actual patterns we observe in the subject in question. The flexibility of thought—the fact that attention to novel evidence would cause them to revise their belief that the clothes are dry or that something is a dead body they would no longer believe that the body they see is that of a dead person. This flexibility demands a different kind of process: a process of inference. And this connection between stimulus-independence and processing in turn supports a connection between stimulus-independence of thoughts and their possession of a representational content. Inferential processes utilize and produce states with representational contents.

Arthur Conan Doyle's Sherlock Holmes. 'A Study in Scarlet,' which documents the memorable first meeting between and Doctor Watson, supplies a nice example of how inferential processes are implicated in capacities to recognize objective states of affairs. Ever the show-off, Holmes announces to Watson his knowledge of the fact that the doctor has just come back from Afghanistan. When Watson asks how Holmes could have known such a thing, Holmes replies as follows:

From long habit the train of thoughts ran so swiftly through my mind that I arrived at the conclusion without being conscious of intermediate steps. There were such

steps, however. The train of reasoning ran, 'Here is a gentleman of the medical type, but with the air of a military man. Clearly an army doctor, then. He has just come from the tropics, for his face is dark, and that is not the natural tint of his skin, for his wrists are fair. He has undergone hardship and sickness, as his haggard face says clearly. His left arm has been injured: He holds it in a stiff and unnatural manner. Where in the tropics could an English army doctor have seen much hardship and got his arm wounded? Clearly in Afghanistan.' The whole train of thought did not occupy a second. I then remarked that you came from Afghanistan, and you were astonished.

We may notice that the form of Dr. Watson's question is essentially the same as one might raise about any systematic behavioral competency: Holmes demonstrates knowledge about a mind-independent state of affairs (Watson's having recently been in Afghanistan), as evinced through a behavioral marker (a linguistic utterance); the question is how Holmes succeeds in getting the answer right given the underdetermination in the evidence available to him. The process by Holmes explains this capacity consists in his recognition of various perceptually available cues which are causal products of the fact in question, and his generating explanations or hypotheses for perceptual data by means of generalizations about relationships between kinds. For instance, he might classify various aspects of Watson's appearance (that his posture is *upright*; that his manner is *reserved*, that his arm is *wounded*) and entertain various hypotheses (a clergyman, someone who has a back brace and is uncomfortable, someone with very strict parents). Holmes draws on capacities to

classify his perceptual presentations as belonging to more abstract types, and generates problems, by noticing inconsistencies, or by being aware of unresolved questions. Holmes notices the color of Watson's skin, together with the color of his wrists, and concludes that this cannot be his natural skin color but must be a tan. The generalization is that since people have the same color of skin all over their body, it can't be his natural skin-tone; and the most likely explanation is that his skin has tanned from sun exposure. As one notices more features, the set of probable explanations shrinks, and one in particular (military personnel) comes into focus. His haggard face shows that he has not been vacationing in the bahamas; since people who have been on holiday usually look relaxed. His arm held in a stiff and unnatural manner indicates injury. Given his profession, this was most likely sustained in combat. Afghanistan is currently the only major war the British are fighting in a hot country. Hence, the conclusion: he has been in Afghanistan. Each of these observations, taken on its own, could be explained by many different reasons or scenarios. But taken in the aggregate, the unique combination of these attributes leads to a single, salient solution.

In conjunction with studying the patterns of a creature's responses, another equally important observation concerns the structure of the underdetermination problem itself. There is a problem of underdetermination whenever a creature's responses to a distal feature of the environment are mediated by their sensitivity to another feature of the environment which is somehow connected to the feature they are interested in, yet which does not determine the presence of that feature. In the case of Ted, the dryer's ending its cycle provides information that the clothes inside it are dry, yet the one fact does not

determine the other. The presence of oleic acid is normally a highly *reliable* indicator of the presence of dead ants. The occurrence of this chemical, independently of dead ants, is extremely rare (the mistakes occur in contrived situations). And, because oleic acid rarely occurs without a dead ant, it is possible for the ants to successfully respond to that property in the environment by means of a very a simple mechanism which exploits that regularity in nature: acid triggers a dead-ant-appropriate response which is normally, fulfils the intended function. The link to the environment can be secured in a purely external way, through functions that 'latch on to' the laws of nature, because the structure of the world itself makes this possible. By contrast, most of the problems of underdetermination we resolve through thought are such that they involve coming up with solutions that are *vastly unconstrained by a given piece of evidence*. And it is partly in relation to this fact that we can understand the need to invoke different kinds of processes and states in explaining those phenomena. In the Holmes example, Watson could have acquired a tan in many different countries, he could have a stiff posture for many different reasons, and so on. Similarly, the perceptual evidence of a lifeless body lying in the street is less regularly correlated with the presence of death than the presence of oleic acid is correlated with the presence of dead ants. Some lifeless bodies encountered in a street may be corpses, but more often they are not (they are more likely to be passed out drunks, or homeless people, for example). For this reason, a system that was programmed to respond in a dead-body-appropriate way whenever they were confronted with that state of affairs would likely get things wrong at least as often as they got things right. All problems of

underdetermination are not created equal—and these inequalities often demand different explanations.

I turn now to the second part of the argument: the link between the stimulus-independence of thoughts and their possession of conceptual content. Once we accept the connection between stimulus-independence and content, it is a relatively short step to the idea that the contents of thoughts must be composed of limitlessly-recombinable constituents, because inference requires conceptually structured contents. As Tyler Burge writes,

At bottom, representational contents are just kinds, or aspects of kinds, of psychological states. The structure of representational contents marks structural aspects of the capacities embodied in the psychological states. The state of believing that the frog has visual perception involves having certain inferential capacities. The belief involves a capacity to infer that something has visual perception, that the frog has perception, that the frog has vision, and so on. These capacities are systematically related to inferential capacities associated with a belief that the frog has auditory perception. These capacities (and the beliefs themselves) have structural aspects inasmuch as they are systematically related to one another and to a more general capacity to carry out deductive inferences. The structural aspects of the representational content of the belief mark structural aspects of the relevant belief, and of inferential capacities constitutively associated with it. (2010: 41)

Just as the notion of a subject's thought about the world becomes meaningful in the context of the subject's possession of inferential capacities in which thoughts essentially figure, so too for the idea of a concept. Concepts show up as a causal explanation for how inference works; and as a principle describing the kinds of representational contents that are produced through this process (the beliefs which are achieved in this way inherit the structure of the states from which they are causally derived. And that process makes essential use of a context-independent representation of a property. A concept of death, for example, is a general understanding of *what it is for a thing to be dead* which is capable of explaining a subject's ability to reason about death in their environment. A recognition of this point also draws our attention to the need to distinguish between the Generality Constraint and the empirical datum which motivates attributions of concepts. The *datum* that concepts are introduced to explain is not systematicity or generality as such but patterns of responses which are explained by capacities for inference. The unrestricted recombinability of concepts shows up in the explanation for these capacities.

The argument I have given in this section provides us with a solid empirical basis for ascribing concepts to a thinker. In contrast to the problems that arise for the argument from systematicity, the argument from inference begins with a secure and non-question begging empirical datum—the patterns of success and error in a creature's capacities to respond to features of their environment—and shows how conceptually-structured contents are directly implicated in the explanation of that datum. The idea of representational contents—and kinds of representational contents—as being licenced by our explanatory purposes has real substance within this story. The notion of a subject's

thought about the world becomes meaningful in the context of explaining a subject's possession of inferential capacities in which thoughts essentially figure. The notion of a subject's conceptual capacities likewise becomes meaningful in the context of explaining that subject's inferential capacities, which rely on conceptual structure. Thoughts are essentially conceptual because *thinking*—the rational process which generate thoughts—makes essential use of concepts in the formation of states with objective content.

3.7 Nonconceptual Content Regained

What emerged from the previous section is a new way of understanding the validity of the Generality Constraint, one which is in accord with Evans's own claim that the Generality Constraint as an essential constraint on thought-contents. Yet not only does the essentially conceptual nature of thought-contents *not* exclude the possibility of perception's having a nonconceptual content, within this framework, it actually gives us a promising basis for believing in the actuality thereof. Whilst the stimulus-independence of perception shows that perception is underwritten by processes which are not information-registration—perception is underwritten by some sort of process which uses and produces states with representational contents—the forms of stimulus dependence that occur in perception also bear on our views about what the nature of those processes might be. As the psychologist, G. Kanisza, points out, it is equally important to consider the differences between seeing and thinking, "because these, by indicating the possibility that the two classes of phenomena obey to different rules, can set us on the road of discovering these rules." Arguing against inferentialist accounts of perception, he writes,

To say that … the visual system 'solves a problem' may be acceptable as a metaphor. Metaphors can of course be useful in sciences as well as in poetry ~ provided that they remain metaphors and are not mistaken for explanations. … (A)ny result of a natural process could be regarded as the solution of a problem. What would be, however, the advantage of such a position? To say that the perception of transparency is the result of an unconscious process of problem solving does not contribute at all to understanding the phenomenon. Our knowledge of the laws determining the phenomenon, of the conditions which facilitate it and of those which hinder or make it impossible, remains as it was before. A metaphor is no substitute for an explanation. Such an 'explanation,' moreover, would have the disadvantage of applying only to positive cases. When a phenomenon does not occur, one can always say that the system has not been able to solve the problem, that it has made some error, that it has let itself be deceived, or applied a rule inappropriately. All this is not a very brilliant way of getting around a difficulty. (1985: 27)

I want to suggest that the limitations or constraints on perception—the forms of stimulus-dependence which motivated our inquiry in the first place—may provide evidence that the processes which underwrite perception are not inferential, and do not make use of conceptually structured contents or conceptual generality.

In thought, there are no a priori limits on the kinds of evidence that could count for a given conclusion. The available routes to the solution of a given problem are in principle open-ended. Suppose your child is not where they are supposed to be: in the kitchen doing their homework. Think of the range of strategies one can employ, in order to discover the information you want to know. You might look for them outside the window (direct perception); you might phone the house of their best friend and ask them (testimony); you might alight on a given hypothesis by going through a list of possibilities through a process of elimination. If one of these routes were to fail to yield results, you could use another. And, if the situation were to become desperate, you might draw on new methods, ones you have never used before. Of course, the available routes will be subject to practical limitations (In practice, a subject may not know that a given kind of datum is relevant to their concerns, or they may not know how to use that information). But things they do not now know are things they could, in principle, come to discover or learn. These things are also aspects of the generality of rational thought. Christopher Peacocke expresses this idea by saying that for methods of discovery in genuinely rational thought, those which are 'canonical' need not be 'exhaustive':

> ... an object's having a certain property is not something constituted by certain methods' having a particular outcome. It is rather something that may potentially be investigated by new means, in need of discovery, new means whose detailed nature cannot be circumscribed in advance.... (O)bjects and properties may be discovered

to have characteristics going beyond the characteristics attributed by application of

the canonical methods. (2001: 614)

Moreover, as he points out, this open-ended character of cognitive methods of investigation

are a necessary prerequisite for the possibility of rational critical thinking. It makes sense

to ask whether certain of one's methods might be inaccurate only if one's methods of

coming to accept contents are ones that aim to give information about states of affairs

whose nature is to be explained independently of such investigation.

Perceptual constancies show that we can arrive at a perceptual 'conclusion' (an

experience of something's shape, or its color, or its size) through different routes; on the

basis of different kinds of sensory 'evidence'. But the range of sensory inputs we are

capable of taking as evidence for a given perceptual conclusion are—again—highly

restricted. You can see a lemon as having a given shape from a range of different angles, or

when it is partially obscured by another figure. But you can't see it as having that shape

through seeing a small portion of its center, or when it is buried out of sight inside a

grocery bag. You can see a lemon as being yellow under a range of different illuminants. But

you can't see it as yellow under conditions of near-darkness—although you could figure

this out through reasoning. There are also dependency relations between certain sorts of

perceptual properties and others. For instance, a perceptual state cannot represent an

object's size without representing how that object is illuminated; nor an object's shape

without representing how the particular is oriented in space. The fact that perception is

subject to these kinds of 'in-principle' limitations suggests a distinctive process underlying

perceptual resolution of underdetermination, one governed by principles quite different to those in thought. The issue we have been focused on so far is that the pattern of responses exhibited by the visual system, in overcoming this problem of underdetermination, is neither like the case of the ants nor like the case of human thought. Perception is more flexible than mere information registration, in that there is no one-to-one correlation between the stimulus and the percept. But the relationship between stimulus and percept is not entirely free and unconstrained as it is in thought: it is not the case that any stimulus could give rise to a given perception; nor that a given stimulus could give rise to any kind of perception. The freedom and flexibility one encounters in perceptual constancies is significantly limited in relation to that found in thought.

These suggestive observations are further solidified when we consider the nature of the underdetermination problem that perception must resolve. Consider a visual perception of a body lying in the street which gives rise to a belief that what is seen is a corpse. The component of the retinal image corresponding to the perception of size is consistent with a range of possible sizes (at different distances) possessed by the distal object that caused that retinal registration. The retinal image projects a shape that is consistent with a range of possible shapes (at different orientations) that figure might possess. The problem here is more difficult than the problem faced by the ants. There is not a reliable regularity between a retinal image and a worldly feature that could ground a purely automatic association between the two. Yet at the same time the underdetermination is less than in it is in thought.

In *The View from Nowhere*, Ernest Nagel notes that the opposition between objectivity and subjectivity is not a binary distinction but rather a matter of degree: "A view or form of thought is more objective than another if it relies less on the specifics of an individual's makeup and position in the world, or on the character of the particular type of creature he is" (1986: 5). The 'proto-rational' and 'proto-conceptual' nature of perceptual capacities may well be regarded as a concrete illustration of this insight. Perception, as a form of nonconceptual representation, rises above the mindlessness of stimulus-dependent registration of information without yet reaching the status of fully-objective status of rational, conceptual thought. The Generality Constraint serves as a theoretical (negative) criterion for nonconceptual content inasmuch as it indicates a 'less objective' form of content. Limitations on a subject's capacities to recombine the constituents of their contents may indicate contents which 'rely more on the specifics of an individual's makeup or position in the world, or on the character of the particular type of creature he is.'

3.8 Conclusion

The guiding idea of this chapter was that we could use the Generality Constraint as a criterion for establishing a nonconceptual content for perception. The two major premises in this argument were: 1) If the Generality Constraint is a valid criterion for a thought's having a conceptual content, then a representational states' failure to satisfy the Generality Constraint is evidence that that state has a nonconceptual content; and 2) Perception fails to satisfy the Generality Constraint. Whilst the tendency has been to take the first premise to be justified by a 'restricted' reading of the Generality Constraint—the assumption being

that it is because contents and concepts are contingently related properties of thought that it is an open question whether there are representational states which fail to satisfy that condition and which therefore have a nonconceptual content—I have shown that the argument that is supposed to justify this interpretation guarantees that conceptuality is a feature of all representational states if it is a feature of any. In its place, I have shown how one can uphold a constitutive connection between concepts and content in thought whilst allowing for the possibility that this connection does not hold for perception. Whilst the arguments I have offered in this chapter for thinking that perception has a nonconceptual content are suggestive rather than decisive, they serve as a clear sign-post for the friend of nonconceptual content. We should not look for nonconceptual content in the space between 'having a content' and 'having a certain kind of content'—for there is no such space to be found. Instead, we should look to the space between information registration and fully objective, inferential thought, as two extremes that explain an organism's engagement with the environment. To move forward from here, what is needed is careful, empirically informed work which meets two conditions: a) identifying that perception involves capacities to respond to features in the environment that require explanation in terms of contents; and b) demonstrating that those same capacities do not require explanation in terms of conceptual contents. This amounts to giving a theory of what it is to perceive: how the nature of the mechanisms underlying perception differ from inferential thought and how they impose different requirements on the structure of the contentful-states involved in these internal computations.

Apart from its implications for the nonconceptualism debate, the discussion of this chapter has consequences for some broad questions in the metaphysics of mind. One concerns the basis for classifying representational states and representational contents as belonging to *kinds*. We presuppose the reality of psychological kinds when we distinguish between a genuine representation of X and merely mechanical operation performing the same function as a representation; or when we distinguish what a person 'really saw' from what they merely thought they saw. Yet there are well-known difficulties in justifying such presuppositions—leading some to question Do they mark genuine joints in nature, or are they merely conventional delineations, reflecting our purposes and interests rather than the nature of reality? With this chapter we have the basis for a robustly realist answer to these questions, one which anchors such distinctions in more fundamental distinctions between the processes which explain successful engagement with environmental entities. States which have representational contents as part of their natures—as opposed to states which are merely described by a representational content —are distinguished by patterns of success and failure in engaging with environmental entities marking which *make use of* and *causally produce* states with representational contents in order to achieve successful responses to the environment. The distinction between kinds of representational states and kinds of contents corresponds to different kinds of processes for which this is true. What it is for something to be a thought about the world is for it to be a state which figures in inferential capacities; and inferential capacities deliver contents and concepts together, in a single stroke. Conceptual capacities are implicated in what it is to have a capacity to represent the world *in thought*, to have a mental representation of that *kind*. The possibility

of nonconceptual content gets a grip, on this picture, in the possibility that there are other kinds of processes, other than propositional inference, through which subjects can achieve representation of the world. What emerges is a meaningful role for concepts and nonconceptual content in psychological explanation.

The present discussion also carries implications for metaphilosophical questions concerning the kind of methodology that is most likely to yield truths concerning the mind. certain reflections on the relationship between a priori theorizing and empirical, scientific investigation and in reaching truths about the nature of mind—particularly in the nonconceptualism debate. This theme goes back to chapter one, where we saw that philosophers of different stripes oppose a nonconceptual content in perception on the basis of a priori theorizing about the essential dependence of representational content on conceptual structure. Whereas chapter two turned a critical lens on this approach, showing where a priori theorizing oversteps its legitimate boundaries, this chapter balances out this picture by criticising a temptation to go too far the other way. Those who maintain that thought is only contingently conceptual—who deny any a priori connection between a thought's function of *representing* the world and its function of conceptualizing the world —do so in part, I believe, because of an assumed identification between what is *contingent* and what is *empirical* or *scientific*. The domain of empirical investigation is a domain of 'open questions,' filled with phenomena whose natures are not circumscribed in advance, which await discovery. The empirical orientation is one I share. But the argument of this chapter shows that the narrative is oversimplified. Those who have claimed the 'scientific highground' for themselves have often ignored the difficult scientific question of when

explanations which appeal to representational contents is genuine. It is because they assume that representation is easy or unproblematic that they think it is in principle possible to have thoughts which function to represent the world without having a conceptual content. When we recognize this as a mistake, what emerges is a more nuanced view of the nature of empirical psychological explanation, in which we see that the contents of thought can be essentially conceptual without precluding the possibility that other kinds of representations may have a different kind of content; and that a priori theorizing is not merely allowable within a scientific approach to the mind, but is integrally part of such an approach.

Considered in relation to the overall argument of the dissertation, the conclusion of this chapter supports and develops the work begun in the previous chapter of establishing the existence of a nonconceptual content in perception. The aim of the previous chapter was to address a conceptual worry about nonconceptual content ('How can we get a grip on the 'very idea' of a nonconceptual content?'). This chapter addressed a kind of epistemic analog of this charge: How could anything count as evidence for the existence of nonconceptual content that was not at the same time evidence that the state in question is not a form of true representation at all? The answers to each of these problems, moreover, connect up in important ways. When classification is viewed as a sophisticated function often identified with the possession of concepts, it looks as though the possibility of nonconceptual content is limited to something that does not organize, structure or classify. But we do not have to think of nonconceptual content in the impoverished way that friends and foes alike have supposed: perception can involve classification without yet rising to the

level of the conceptual. By turning a critical eye on existing arguments for nonconceptual content in this chapter, we have discovered—and discounted—a corresponding prejudice concerning what would count as evidence for nonconceptual content. We do not have to deny—per impossibile—the existence of capacities to recombine the constituents of perceptual content for perception to have a nonconceptual content. Perception can involve generality without yet rising to the level of conceptual generality. Now that these barriers to nonconceptual content have been dispelled, the work that remains is to leverage this suggestive framework to reach a firm conclusion.

4. THE PROBLEM OF PERCEPTUAL RECEPTIVITY

4.1 Introduction

Perceptual states inform of the presence and nature of objects in our immediate environment, based on input to the senses. Yet we intuitively distinguish perception from other kinds of states, such as judgements and beliefs, which may also be said to inform us about the presence and nature of objects based on input to the senses. What must be added to our concept of perception to distinguish a visual experience of a yellow lemon on the kitchen countertop from inferring the presence of a yellow lemon on the kitchen countertop based on evidence or the testimony of others? Whilst this question—what distinguishes perception from cognition?—is different from the question we started from—does perception have a nonconceptual content?—the previous chapter has showed us that the two are intimately connected. A theory of perceptual content is answerable to a theory of the processes which produce perceptual content.

The idea I shall explore in this chapter is that the mark of perception lies in *the way* that the senses give rise to objective perceptions: perception is a 'receiving of the world.' The image of receptivity is due to Kant, for whom it forms the basis of sensible representation. In the Critique of Pure Reason, Kant evokes the contrast between sensible and intellectual representation in the following terms:

Let us give the name sensibility to our mind's receptivity, [i.e., to its ability] to receive' presentations insofar as it is affected in some manner. Understanding, on

the other hand, is our ability to produce presentations ourselves, i.e., our spontaneity of cognition. Our intuition, by our very nature, can never be other than sensible intuition; i.e., it contains only the way in which we are affected by objects. (B75/A51)

The Kantian notion of 'Intuition' is generally taken to correspond to what we would now call a 'perceptual experience'; and 'Sensibility', the faculty which generates Intuitions, to perceptual capacities (vision,   audition,   and so forth). Kant's claim is that what distinguishes perception from conceptual thought is the kind of process by which they are formed. Judgements arise through a process in which the subject themselves plays an active role in constructing the information they possess: subjects 'produce presentations themselves' from out of their perceptions. Perceptions are different kinds of states from judgements in part because they are not things a subject constructs out of something else in this way; instead, perceptions are simply 'given' to subjects who are appropriately equipped simply through coming into sensory contact with the object. The same idea has been conveyed by contemporary philosophers through the image of 'openness.' For example:

In experience one finds oneself saddled with content. … The content is not something one has put together oneself, as when one decides what to say about something … (I)n enjoying an experience one is open to manifest facts, facts that obtain anyway and impress themselves on one's sensibility. (McDowell, 1994: 29).

(P)rimitive objectivity does not depend on individuals' producing it. Individuals do

not construct objective perception from subjective representation or consciousness.

... [Perception] starts with an openness to the physical environment as it is. (Burge,

2010: 547)

At its core, these images of 'receiving' or 'openness' speak to an intuition we have about the

source of perception, what is bearing the bulk of the burden in explaining our possession of

objective information. Thoughts, being constructed out of *experiences*, come partly from the

world. But in being *constructed* out of experiences, they also come from the subject: to

count as genuine thought it must be the subject who is shouldering the bulk of the

representational burden (Camp, 2009: 288). In speaking of perception as a receptivity

rather than a construction— something given rather than achieved; something we receive

rather than something we do for ourselves—what we are really saying is that in genuine

perception, it is the world and not the subject that bears the bulk of the representational

burden. Perception is that awareness of the world which we get *from the world*. The flow of

information is 'world to mind' rather than 'mind to world.'

The problem of perceptual receptivity I reference in the title of this chapter arises

from a tension between two plausible assumptions. On the one hand, it is natural to

suppose that for the receptivity intuition to be vindicated, the objective information that

perception presents must be 'contained in' the sensory inputs to perception. For if we

invoke a process of 'construction' in the formation of judgements, this is to explain the fact

that the information we grasp in judgement is *not* contained in our perceptual experience: judgements take a stand on a question that is not closed by the perceptual experience, so it falls to the subject to settle the question themselves. On the other hand, we regard the inputs to perception as sensory registrations; and the problem of underdetermination in perception reveals that the immediate registrations of objects by our retinas don't close the question for perception, either. Thus, whilst there seems to be something importantly right in the idea that we perceive only what the world 'gives' us, there is also a reasonable worry that this could not possibly be right. I believe the receptivity intuition deserves to be taken seriously. But an account of receptivity must also take seriously the theoretical problem of underdetermination that seems to threaten it. That is, we can't abandon the receptivity intuition in our attempts to accommodate the scientific data, as cognitivist or constructivist theories of perception attempt to do; but nor can we abandon the scientific data in our attempts to do justice to the receptivity intuition, as naive realists have sought to do.

I believe that Kant's account of perception contains just the tools we need to resolve this problem. The first step is to recognize that Kant is offering us a way of understanding what it means to speak of receptivity in relation to objective perception, as distinct from merely subjective sensation. In the Transcendental Aesthetic, Kant writes,

Now, the subject's receptivity for being affected by objects precedes necessarily all intuitions of these objects. Thus we can understand how the form of all appearances can be given in the mind prior to all actual perceptions, and hence given a priori; and we can understand how this form, as a pure intuition in which all objects must

216

be determined, can contain, prior to all experience, principles for the relations among these objects. (A26/B42)

Whilst many interpreters of Kant have read 'receptivity' as a purely passive operation, I argue that passages like this one show that Kant is in fact offering us a way of reconceiving perceptual 'givenness' in a way that does not exclude the mind's internal contributions in the creation of meaning. The contrast between receptivity and spontaneity, as means of producing objective representations, lies in a notion of sequence or order. The process of thought starts with a perceptual presentation of an object and proceeds to generate knowledge pertaining to that object by our 'scrutinizing appearances to find some rule in them.' The organizing or interpretative operation contributed by perception, by contrast, is not something *performed on* a presentation already given to generate something new. Rather, it is one in which what organization the mind contributes *precedes* what is organized. We have internalized a fixed and unitary 'system' of spatial categories and the laws of euclidean geometry which specify the laws relating those categories to one another; and perception is a matter of fitting sensory registrations 'in' to that pre-existing framework. I support Kant's suggestive remarks by showing how they find a concrete realization in recent predictive processing theories of perception, which characterize the perceptual process as one of generating predictions or hypotheses, based on Bayesian updating, about what the sensory input from a previous processing stage will be.

The second step is to recognize that this process yields a kind of 'nonconceptual content' which preserves the core of the 'containment' intuition. The opposing forms of

attribution I argued for in chapter two emerge as the consequences of this more fundamental opposition between processes for generating objective representation: "Concepts, therefore, are grounded on the spontaneity of thought, as sensory intuitions are on the receptivity to impressions" (A68/B93). Spatial attribution allows us to see that, whilst the objective content of perception is not *determined* by sensory registrations, we preserve a notion of containment by recognizing that the sensory component of perception partly constitutes—and thereby constrains—the objective content. Whilst this proposal is revisionary with respect to some of our naive intuitions—notably in allowing for a notion of internal processing and classification within perception—I think it would be a mistake to think that a satisfactory solution to the problem of perceptual receptivity must be one that preserves our naive intuitions in their entirety. Although we may not get unalloyed receptivity, in the end what we get is good enough—good enough to give us a robust criterion for distinguishing perception from thought and to yield a form of objectivity to stand   alongside   rational,   conceptual   thought.

4.2 The Intuition

Consider the following scenarios:

> *Lemon 1*. Imagine you are standing before a barrel piled high with lemons. One of
> these is covered by other lemons in such a way that its edges are obscured and only
> the central part of its surface is in your line of vision. Directing your attention on this
> lemon, you see the central portion of the surface of the lemon, including its bright

yellow color and the dimpled, shiny texture of its skin. Recognizing this distinctive color and texture as proprietary to lemons, you conclude that the object you are seeing has a lemon-like shape.

Lemon 2. Imagine there is a lemon on the kitchen counter, and the kitchen is dark save for a chink of light coming from a streetlamp outside. Looking at the counter, you are able to discern the distinctively lemon-like outline of an object, with its rounded central body and nobbly ends. Recognizing this outline as the distinctive shape proprietary to lemons, you conclude that the object you are seeing is yellow.

As these examples illustrate, we have clear-cut intuitions that there is a difference in kind between genuinely perceiving an X and making a judgement or drawing a conclusion about an X based on perceiving something else. In *Lemon 1*, we have the definite intuition that your visual experience of the lemon in the barrel includes a perceptual presentation of the color and texture of its surface and excludes its shape. Although you are in possession of information about the lemon's shape, this information takes the form of a judgement you make based on the perceptual presentation of the color and shape of the lemon, and is not itself perceived. In *Lemon 2*, we have the definite intuition that your visual experience of the lemon under near-dark conditions presents the lemon's shape but does not include a presentation of its surface color. Although you are in possession of correct information about that its color, this is a judgement you make based on the shape you perceive, and not itself something that is perceptually presented. An important thing to notice about these

examples is that all of the information you possess about the lemon can be accurately said to be acquired 'through the senses.' For example, assuming that the lemon that you see actually is yellow and possessed of a shiny dimpled texture, you are aware of these properties possessed by the mind-independent lemon by virtue of some causal interaction between the lemon and your retina. Yet it is just as much true to say that you are informed of the lemon's shape based on that object causally affecting your eyes—since you make this judgement on the basis of what is presented to you in perception.

A natural way to motivate the receptivity intuition is to point out that it supplies us with an explanation for our rather definite intuitions about these cases. The receptivity intuition is that our ordinary concept of perception includes the way in which information is acquired through the senses: a subject perceives an X if and only if they have an awareness of X that is 'received through the senses' or 'given by the senses.' Our intuitive classifications of states of awareness into perceptions and non-perceptual judgements tracks this criterion about how we come to be in possession of this information.

Consider *Lemon 1*. Suppose we ask: how it is that you come to be aware that there is an object in the barrel with that color and texture? A natural answer is that the color and texture of the lemon are simply 'present' or 'there' in your experience. It as though you can mentally 'point' to those things. Since these properties are already present or given in sensory experience, they require no processing other than transmission to become known. There is nothing you have to *do* anything to acquire this information about the lemon. We think that what is taken in through the senses is sufficient to explain the formation of the perceptions. The objects of your perceptual experience are 'given' to you in an immediate

220

way, simply by being equipped with a visual system and having these things within your line of sight. Now suppose we ask: how is it that you come to be aware that there is a lemon in the barrel with that shape? The same answer is not available, because the shape is not 'present' or 'contained in' your sensory experience. One way to bring out this idea is to point out that you might believe that the object has a lemon-like shape whilst being unsure of the exact proportion the nibs at each end stand in to the circular body, or of the size the lemon has overall. Since it is not contained in your sensory experience, there must be something you do which explains how you come to be in possession of this information. Perhaps you inferred it based on your recognition of the color and texture of the surface, together with your background knowledge that most objects with that color and texture are lemons. Or perhaps you inferred it based on the observation that all of the surrounding items have that shape; hence the statistical likelihood is that the obscured item also has that shape. The main point is that, whatever the correct answer turns out to be, there must have been some kind of constructive operation which explains how you got from the sensory experience to the information, because without that contribution the sensory experience isn't sufficient to account for your possession of that information. have that information if I didn't do what I needed to to get it. The intuition that your awareness of the lemon's shape is a judgement *rather* than a perception seems to be tied to this idea that the way in which you come to be in possession of that information concerning the lemon is different. The reason we are inclined to classify certain states as judgements *rather than* perceptions is because they are constructed out of what is received through the senses— and this is not the right kind of process to count them as perceptual.

A further way to support the receptivity intuition is to consider the conditions under which we tend to revise our classifications of mental states from perceptions to thoughts. If we encounter evidence that a putative perception is sensitive to contextual cues this tends to cause us to revise our classification. For example, consider the following:

Many victims thought [the killer's tears were real]. But there is an exception amongst the victims who believed the tears were real. "She heard those sounds, but she wouldn't attribute them to crying". What did she think it was? "High pitched hysteria. Like laughter.

The description that is given here would not under ordinary circumstances be the most natural one: instead of "I heard a sound and thought it was crying" we would more often say "I heard (them) crying," or "I heard laughter." What makes it natural in this case is that, since different subjects heard the same sound in different ways, 'what is heard' is subject to *interpretation*. The crying itself can't have been 'perceptually given' to them; rather, they must have played a role in constructing it based on something else they *did* hear. Genuine perception is immune to influence from our beliefs in this way. Suppose I know than an experience I am having in which there appears to be a yellow lemon before me is in fact illusory: the object has a white surface onto which a yellow light is being directly shined. My knowing this won't change the fact that it *appears* yellow. As Crane and French point out, this idea that a genuinely perceptual presence *imposes* itself on the observer, without

being subject to influence by other mental states, is a natural consequence of the view that perception is receptive rather than constructive.

> in the case of thought, how the object of thought is at the moment one is thinking of it does not in any way constrain one's thinking of it; but in the case of perception it does. (I)n the middle of winter, one can imagine the churchyard as it is in spring, covered in autumn leaves, and one can think of it in all sorts of ways which are not the ways it presently is. This is not available in perception, because perception can only confront what is presently given: in this sense, it seems that you can only see or hear or touch what is there. (2017)

The reason why we can't influence what we see or hear is because we see or hear only what is *there*. If perception is a receiving of the world, we might say, then in perception you 'get what you are given.' By contrast, that information we have about objects which we play some role in creating—by constructing it out of what we are given in perception together with background assumptions—is also subject to change. If the assumptions we bring to bear are partly responsible for our way of taking the world, then differences in those assumptions may give rise to differences in how we take the world to be. This freedom that is present in thought is a key reason why it makes sense to hold thinkers responsible for the conclusions they draw: it only makes sense to blame someone for something if they had some choice in the matter.

The idea of what is present in experience, which plays a central role in the receptivity intuition, seems connected to the idea of phenomenal character in ways which it is important to clarify. The phenomenal character of an experience of X is the conscious, subjective, qualitative aspect of an experience ('what it is like' to perceive a shiny yellow lemon) as opposed with its objective 'content' (the mind-independent facts that are conveyed by an experience). (Note that I am using the term 'content' here to signify the objective significance of the experience, whilst leaving open whether that objective significance is constitutively an abstract object determining a set of correctness conditions.) Those aspects of my awareness that are 'contained in' my sensory experience—for example, the visible portion of the surface of the lemon, together with its yellowness and its shiny and dimpled texture—are things I'm aware of, in a particularly vivid or immediate way. The shape of the lemon that is not contained in my experience lacks this vividness or immediacy. But although this vividness is experiential or subjective aspect of my experience, I think it would be a mistake to equate this with phenomenal character, since the two can come apart. an example Wittgenstein gives, in 'Remarks on Color', of an experience he has when looking a black-and-white photograph. He describes this experience as follows:

> I see in a photograph (not a colour photograph) a man with dark hair and a boy with slicked-back blond hair standing in front of a kind of lathe, which is made in part of castings painted black, and in part of smooth axles, gears, etc., and next to it a grating made of light galvanized wire. I see the finished iron surfaces as iron-

coloured, the boy's hair as blond, the grating as zinc-coloured, despite the fact that everything is depicted in lighter and darker tones of the photographic paper.

He then proceeds to frame a question:

64. But do I really see the hair blond in the photograph? And what can be said in favor of this? What reaction of the viewer is supposed to show that he sees the hair blond, and doesn't just conclude from the shades of the photograph that it is blond?- -If I were asked to describe the photograph I would do so in the most direct manner with these words. If this way of describing it won't do, then I would have to start looking for another.

Wittgenstein's description of his experience at face-value, he has a visual *experience* which presents the color, blond, as immediately and non-inferentially present in the photograph (specifically, in the part which represents the boy's hair). Nevertheless, the intuitive thing to say is that he does not 'really see' the hair as blonde. The blondeness is a construction or interpretation of what is immediately given to our eyes: he reached that conclusion based on seeing the black-and-white image together with prior familiarity about the objects depicted and how they would look in a black-and-white photograph. For there is another, more important sense in which the blondeness clearly is *not* 'present in' the photograph: the property, blonde, is a color; but the photograph is by stipulation achromatic, *colorless.* This suggests that the notion of presence that is relevant to the receptivity intuition is

distinct from the more general notion of perceptual presence. The 'immediate presence' of the blondeness is only an *experience* of immediate presence; the mind's projection onto something which does not actually contain it. (A comment Wittgenstein adds later makes this point even more obvious: "If the word "blond" itself can sound blond, then it's even easier for photographed hair to look blonde!") This is distinct from some other, 'informational' notion. Perceptual presence in the second sense is a necessary condition for a state to count as genuinely perceptual.

Another reason to think that there is a difference between the notion of 'perceptual presence' that is relevant to the receptivity intuition and the ordinary notion of phenomenal presence can be appreciated by reflecting on perceptual illusions such as the phenomenon of illusory contours, as illustrated by the Kanisza triangle:
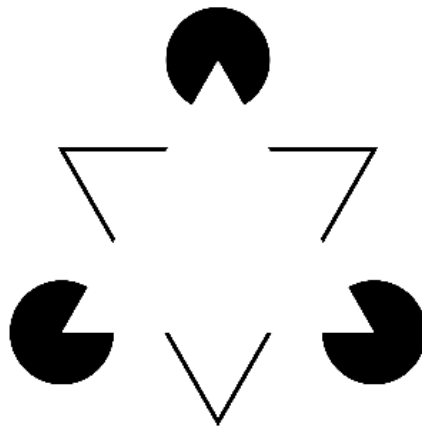
Figure 3. The Kanisza triangle demonstrates the phenomenon of illusory contour. The central figure is seen as a triangle even though no edges are marked on the page.

In this example, we have a definite experience of a triangle, with three corners and three edges, as being 'there' at the center of this image. This is one thing that makes us inclined to classify this experience as a *perceptual* phenomenon rather than a cognitive one (compare the case of the underlying image in which what you see are only the corners of the triangle but not the edges: although you might surmise that this is a triangle with its edges obscured, you do not see them as present). Yet in another sense we know that the triangle is not 'there.' I cannot point to the part of the image that corresponds to my perception of the edges of the triangle, in the way that I can point to the part of the image that corresponds to my perception of the three corners of the triangle or the three incomplete circles. What's uncanny and unnerving about illusions such as these is that they reveal, through a weird sort of dissonance, that there are two notions of presence—which normally go together yet which may go apart. When we say that the objects of our experience are 'there' and so are simply received or taken in by us, what we take ourselves to be saying is that they are there outside of us, in the world. The right sense of receptivity takes in what's there *in the world*; the other term of the relation is *the world*. But what seems to emerge is that there are two notions of 'there' there: 'there in our experience' and 'there in the world.' Although the triangle seems to be there, it is not *really* there—or better, it is there in the wrong kind of way. For if it were the right kind of presence, it wouldn't be an illusion; I would be experiencing what is there. There are two notions of presence, and the one associated with receptivity is more basic than, and in some sense explanatory of, the phenomenal one.

The notion of presence I am concerned with plays a central role in arguments for naive realism. The naive realist rejects a representational account of perception in favor of a view on which perceptual experiences are direct relations between perceivers and constituents of the mind-independent world. One of the main motivations for this claim (if not the main motivation) is the need to capture the phenomenal character of experience: "The phenomenal character of your experience, as you look around the room, is constituted by the actual layout of the room itself: which particular objects are there, their intrinsic properties, such as color and shape, and how they are arranged in relation to one another and to you" (Campbell, 2002: 116). When naive realists say that phenomenal character is *constituted by* the things out there they are reacting against the tendency in contemporary philosophy to talk about phenomenal properties as though they are mental properties individuated by what it is like to experience them. One of the important consequences of this discovery is that the naive realist's claim turns out to rest (in a way that has not been widely acknowledged) on the receptivity intuition. They want to say that we are in direct contact in the world—because this gives us the right way of capturing the nature of our experience. But in virtue of what do these facts about perceptual experiences hold? Why does perception consist in a direct contact with the world? Because the way it is brought about is through a receiving of the world. As A.D. Smith puts it,

> Perceptual consciousness is, at least when veridical, an immediate registration of a normal physical object, in the sense that the sensory character of your conscious state… is accounted for by the possession by that object of perceptible qualities,

together with the fact that you stand in a relation of awareness, or receptivity, to it"

(Smith, 2002: 43-44).

This claim is also echoed by Crane and French, who note that it is *because* perception is a

taking in of what is there that perception has an immediacy or vividness which thought

lacks: "this vividness derives from the fact that perceived objects and their properties are

actually given to the perceiver when being perceived, and determine the nature of the

character of the experience" (Crane and French, 2017). The central naive realist claim is

about the structure of perceptual experiences: a direct relation between perceivers and the

constituents of the mind-independent world. The reason why experience has this structure

is because of the way in which you come to have a perceptual experience of objects—by

receiving them.

4.3 The Problem

I have tried to show how the receptivity intuition is central to our commonsense concept of

perception. Yet there is a question about how this intuition could possibly be true.

We started out with the idea that judgements are constructed by a subject out of

experience because they are not 'contained in' sensory experience. What do we mean by

this? One way to understand the claim that a judgement to the effect that a is *F* is not

'contained in' an experience of *a* is in terms of the idea that the experience does not

determine the *content* of the judgement. The experience rationally supports more than one

content that could reasonably be inferred from it. In *Lemon 1*, your experience of the

surface of the lemon is consistent with the possibility that the object you see has a lemon-like shape; but it's equally consistent with the possibility that that object is spherical, or an irregular cube with a curved surface, or that it is simply a curved surface with nothing behind it. In Lemon 2, your experience of the lemon on the kitchen counter is consistent with the possibility that it's a yellow object in dark light. But it's equally consistent with the possibility that its a green object, or a blue one, or a grey one. Since the experience is equally consistent with other states of affairs, we need some explanation for why the subject selects this alight on one possibility rather than any of a variety of other possibilities (particularly in cases where the possibility they choose is the *right* one).There's an open question just from your experience what color the lemons are. The question, then, is how the question does get closed: how do you know what the shape of the lemon is, given that the experience is silent? The image of construction or producing arises as an answer to this question; it offers a way of closing the epistemic gap. Closing the question is something *you have to do* because the senses don't close it by themselves. You draw on background knowledge which singles out one of these possibilities as *more likely*. Notice that it's not just that you are performing an action, it's that you have to supply resources in the form of information—knowledge about how the world works—because the experience itself is *impoverished*.

These criteria by which we determine that thought is constructed out of experience are also seemingly present in perception. The classical puzzle of perceptual theory, 'the underdetermination problem,' is to explain how we perceive the world accurately (at least for the most part), given that there is nothing in the structure of the immediate stimulus

which is specific to its source. The inputs to perception are sensory registrations: a visual experience of an object's size is formed on the basis of the extension of an edge on the retina; a perception of shape is formed on the basis of the two-dimensional registration of patterns of light on the retina. But the same image on the retina, for example, would seem to be consistent with an infinite set of possible circumstances in the world. An illustration of this which is a close analog to the shape example is the phenomenon of amodal completion.
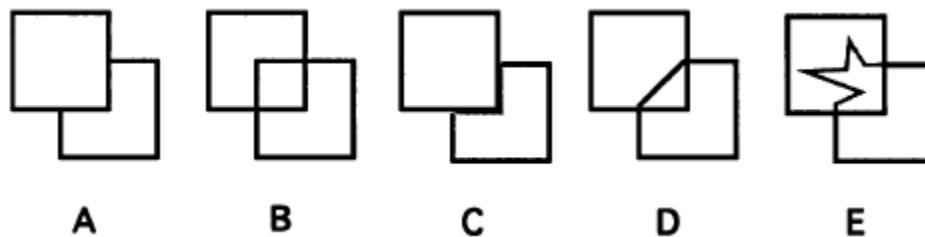
Figure 4. An illustration of the visual phenomenon of amodal completion. The visual system 'fills in' information about the shape of the obscured object, choosing the 'solution' in figure B rather than any of those in figures C-E. The illustration comes from Palmer (1997: xiv).

The pattern of fig. A is usually seen as one square partially covered by another, as in figure B. Yet notice that the image doesn't actually contain two squares. There are two regions, one of which is a square, the other of which is L-shaped, as shown in figure C. The pattern of figure A is just as much consistent with figures D or E as it is with the possibility described in figure B. This example serves as an illustration for a phenomenon that is widespread in ordinary perception. If I look around the room I am in at present, many of the objects I perceive are not in their entirety (and the back sides of objects are all obscured), yet this does not prevent me from seeing the surface of the desk as rectangular, the lamp shade as cubic, or my coffee mug as cylindrical. We can take this idea a step

further, and point out that even the items which are unobscured are not determined by the information that is registered on my retina. For example, as I look at the cylindrical coffee mug before me, what is registered by my retina is only a two-dimensional configuration of lines and an ellipse. This image is consistent with a range of possibilities concerning the actual size, shape and orientation of the mug. This derives from the fact that the retina, which can register information from only one point on any given line of sight, is a two-dimensional array which does not contain information about depth. Without depth information, it is also impossible to determine the spatial properties of an object: its length or the width or shape.

The perception of color involves a similar kind of underdetermination. Consider the two images below:
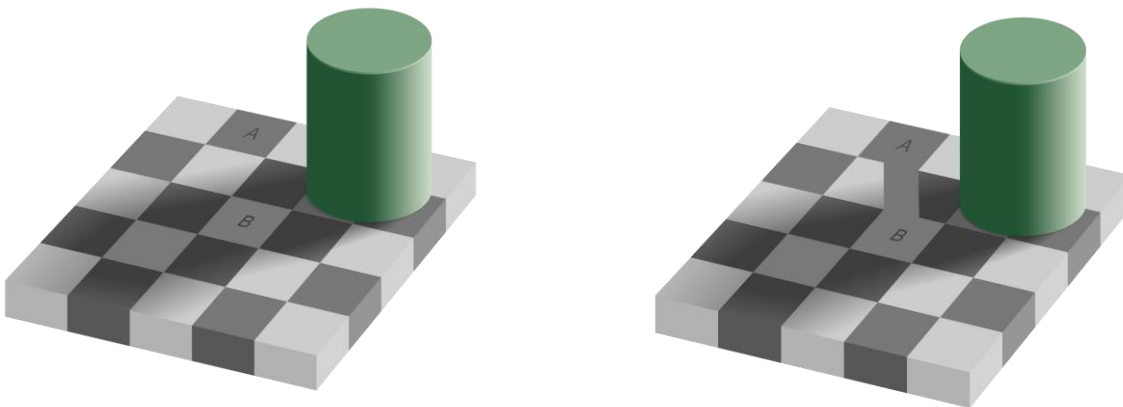
Figure 5. The chessboard illusion illustrates the processes that underwrite normal cases of color constancy. In the left figure, squares A and B are seen as different colors due to visual cues indicating that B is in shadow. In reality, the two squares are the same color, as illustrated in the figure on the right.

In looking at the first image, one has a definite experience of the two squares, A and B, as being of different brightnesses: A is experienced as dark gray, B is experienced as light gray. But notice that the squares are actually the same lightness, as the second image demonstrates. Color is a property of the surface of objects, corresponding to the intensity and wavelength of light reflected by the surface of objects (the 'distal stimulus'). We are able to perceive the colors of objects because our retina is sensitive to the intensity and wavelengths of light striking the retina (the 'proximal stimulus'). But the proximal stimulus confounds the light reflected by objects with the wavelength and intensity of light (the intensity of light registered by the retina is the joint product of surface lightness and how surfaces are illuminated). For this reason, a registration of a given intensity and wavelength is consistent with a variety of different causes: a black surface under bright light, a gray surface under average-intensity light, a white surface in dim light. In the chessboard illusion, the visual system 'reads' the ambiguous stimulus marked 'B' as a light square in shadow because the presence of the cylinder casting a shadow over the board provides cues about shadows; whilst it reads the ambiguous stimulus marked 'A' as a dark square because there are no such cues. Again, this serves as an illustration for the problem faced by the visual system in determining lightness and color constancy in normal circumstances: the inputs to perception underdetermine the contents of perceptual experiences formed on their basis in much the same way as experiences underdetermine the contents of judgements.

A second thing we might mean by saying that judgements are not 'contained in' the experiences on which they are based is that the experiences underdetermine the

judgements themselves. A subject's judgements about what they are seeing can change even though there is no change in their visual experience. In *Lemon 1*, my experience of the texture and color of lemon in the barrel causes me to form the judgement that the object I am seeing has a lemon-like shape. But the experience does not *necessitate* that judgement: if new information were to come to light—if I were to learn of another kind of (differently shaped) fruit that has the same color and texture of skin, for example—I might revise or withhold my judgement concerning the shape of that item. This claim about judgements also has an analog for perception. It is true, as we have seen, unlike a merely imagined or 'thought' phenomenon perception is not easily subject to influences from our beliefs, desires or intentions. But perception may be subject to influence by other *perceptions*. Consider, as a well-known illustration of this, the McGurk effect which demonstrates cross-modal influence in the processing of speech phonemes. The effect occurs when the auditory component of one sound (e.g. an auditory perception of the phoneme 'BA') is paired with the visual component of another sound (e.g. a visual perception of a mouth uttering the phoneme 'GA') leading to an alteration in the perception of the original sound ('BA' is now heard as 'GA'). The sound does not change. Yet your visual perception of mouth movement makes it sound completely different (when you close your eyes and listen to the sound the effect is immediately cancelled). In a BBC documentary which studies the McGurk effect, the presenter takes this as evidence that the perception of speech phonemes is not truly perceptual. that "the mouth movements we see when we look at a face can actually influence *what we believe we are hearing*. If we close our eyes, *we actually hear the sound as it is*… Remember, the only sound you are hearing is 'BA' with a 'b.'" This is based on the idea

that, if perception were subject to influence based on other perceptions without any apparent change in the sensory stimulus, that would be evidence that the subject is partly responsible for bringing about their perception. The problem is that we shall then seemingly have to allow that very few, if any, of the phenomena we intuitively classify as perceptual will turn out to be genuinely so. The last few decades of research in vision science contains a wealth of experimental evidence demonstrating the importance of relations among percepts in the production of an organized perceptual world. Gogel and Koslow (1972) found that if a stationary point of light is observed in a luminous frame in otherwise dark surroundings, setting the frame in motion will induce the stationary point to appear to move in a direction opposite to the physical movement of the frame. Thus, "(o)ne perception, the perception of the depth relation between the frame and the point of light affected another perception, the perception of the motion of the point of light" (ibid., 213). Rock and Brosgole (1964) demonstrated that the perception of 3-D proximity between the dots in an array determines the perception of grouping: tilting the array does not change the perceived organization as long as observers have sufficient information to perceive depth accurately; moreover, if the relative spacing of the objects is misperceived due to lack of sufficient depth information (e.g., viewing the array with just one eye) or due to the breakdown of depth perception under extreme conditions (e.g., large angles of tilt), perceived grouping does switch to rows. Gilchrist (1977) showed that perceived lightness is affected by perceived spatial arrangement, by manipulating depth information so as to cause a target to appear coplanar with either one or the other of two surroundings of objectively different luminance. If "perceived lightness of a surface can vary from white to

black depending merely on its perceived spatial position, without any significant change in the retinal array," Gilchrist argues, "depth processing must occur first and be followed by the determination of surface lightness... the nervous system uses various depth cues to construct a spatial model to fit the retinal pattern" (ibid., 186-187). (See Epstein (1982) for a detailed overview of the experimental evidence for percept-percept coupling).

The problem of underdetermination poses a problem for the receptivity intuition. We regard the process which gives rise to judgements as a constructive process because the judgemental outputs of that process are not 'contained in' the perceptual inputs, and so need to be worked up out of that input by some sort of generative process. We think that the process leading from the senses to objective perception is a receiving *rather than a construction* out of something else more basic because the objects of perceptual awareness are 'there' or 'immediately present' to us. Rendered in more formally precise terms, the claim is that the information we are aware of in perceptual experience is contained in the sensory inputs to perception. That is,

> Perception is a receptivity to aspects of the mind-independent world if and only the objects of perceptual experiences are already 'contained in' the sensory inputs to the perceptual process, so that they do not require any process other than transmission or attention to become known.

The problem is that once this is granted, perception fails to meet the criteria that would qualify it as a receptivity. The information that is immediately registered by the senses

does not 'contain' the perceptions formed on their basis, on any plausible way of understanding 'containment.' The contents of our perceptual representations are not and cannot be determined, one-one, by the sensory array that makes first causal contact at the site of proximal stimulation. And the perceptual states which arise on the basis of sensory registrations are not determined by their immediate sensory causes, either.

One response to these observations would be to abandon the receptivity intuition. Many theorists of perception within the cognitive sciences have taken the underdetermination problem as evidence that "perceptions are constructed, by complex brain processes, from fleeting fragmentary scraps of data signalled by the senses and drawn from the brain's memory banks themselves snippets from the past" (Gregory, 1974: xviii). The leading constructivist account regards perception as a 'ratiomorphic' problem-solving process, in which the visual system engages in a kind of inference to the best explanation for what caused the ambiguous stimulus-array, similar to the way in which scientists make judgements about perceptual data based on hypotheses about their likely causes. Although "the dependence of perception on sensory information makes for certain differences between it and "higher" cognitive functions such as imagination and thinking," as Irvin Rock expresses this idea, "perception is intelligent in that it is based on the same kinds of operations that characterize thought… the processes underlying these different mental end products may be quite similar…. perception results from cognitive operations" (Rock, 1982: 525). According to these philosophers, then, there is no 'problem' about perceptual receptivity because perceptual receptivity does not exist. They are happy to abandon the idea that perception has its source in the world in any sense that is not also

and equally true for empirical thought. In the words of J.J. Gibson, a famous critic of cognitivism, "the activity of perception is supposed to be an internal or subjective process. Meaning is supposed to come from inside, not from outside" (1967: 64). Yet I do not think the receptivity intuition can be so easily cast aside.

Let us consider more closely Rock's suggestion that the main difference between the perceptual problem-solving process and genuinely cognitive inference is that perception is constrained by what is present in the stimulus. The dependence of perception on the stimulus, which Rock himself acknowledges in the passage quoted above, generates the very same tension regarding receptivity within the cognitivist view. Rock discusses two sorts of constraints on perception imposed by the stimulus (1983; 1985). First, the perceptual solution to the problem of underdetermination must be 'supported' by the stimulus, in the sense that "the stimulus must contain the essential features that are called for by the internal solution, i.e., those features that would be expected to be present in the stimulus were the outer object or event were such and such." For example:

> (I)f informed that the object is a cube, then the stimulus must contain the outer and inner contours of a cubic or, to take another example, if parts of the contours of an outline triangle are missing ... then while one can agree that a triangle might be what the contours represent, one cannot perceive a triangle. If, however, a rationale is provided for the absence of these contours, as when another figure, real or illusory, is present that would partially occlude the triangle, then this would be stimulus support for the perception of the triangle. ... It seems that the perceptual

system must 'see' for itself, as it were. those features that must be present were the hypothesized object or event to be present. Knowing that those features are present is not good enough. (1985: 11)

Although I think the term 'stimulus support' is unsatisfying as a description of this phenomenon for reasons we shall consider shortly, the difference between perception and thought which Rock is evoking here seems clear enough. When you have a belief about the shape of the partially visible lemon, that belief is in some sense supported by your experience: you experience an object which you know must have a shape, and that is the most likely shape given the perceptually available data. But you can also have beliefs about my experience which do not have any basis in the experience itself; for example, you can believe that behind the visible lemon is another lemon-shaped object which is completely concealed from view, or that there is someone standing beside the barrel wearing an invisibility cloak. By contrast, given what is immediately presented to you, you couldn't *perceive* these things because there is no stimulus support for these logical possibilities. We might also apply this idea to color constancy: there has to be stimulus support for an interpretation of color; if there's no 'color' in the stimulus then it's not possible to perceive the color that way (Wittgenstein's example of the black-and-white photograph is an illustration of this latter idea). Alongside this, Rock identifies a second kind of constraint: the perceptual solution to the problem of underdetermination must 'conform' to the stimulus, in the sense of being 'compatible with' the pattern of retinal stimulation. For example, one cannot perceive a stick that looks bent in water as straight because that is

inconsistent with the retinal image—although one can believe that it is in fact straight (the appearances are illusory). Given the sensory experience you have of the lemon in the barrel, you cannot perceive there as being no object present at the region where you direct your attention—although you could believe there is no object present (for example, if you believed yourself to be hallucinating).

Whilst Rock tries to explain these stimulus-constraints within an inferentialist view of perception, the difficulty is that their descriptions implicitly presuppose the receptivity intuition. Consider the claim that the solutions available to the perceptual system must 'support' the perceptual possibility. What exactly do we mean by this? Judgements, too, must be supported by evidence of some kind. Under what conditions can it be said that the *stimulus* supports the 'perceptual conclusion'? How shall we analyze the supporting relation so as to make clear the difference between perception and thought? When Rock says, 'It seems that the visual system must "see for itself" the features that would be expected to be present, he uses a telling metaphor: 'perception must 'see for itself' enough of the triangle to support that solution' looks awfully 'there has to be enough of the triangle 'in' the retinal image for perception to interpret it as a triangle' (or again: 'there has to be enough color 'in' the stimulus for perception to interpret it as colored'). One observation that seems relevant here is that in perceptual constancy, the proximal stimulus *confounds* two sorts of information: information about objects and information about a subject's perspective. Perception can disambiguate an ambiguous stimuli; but it can't conjure possibilities (even nomologically possible ones) on the basis of nothing. So even if there is some amount of building that takes place, there is some information that must be received

240

even at first sensory contact. These 'constructive' processes cannot get started unless there is something *in the stimulus* that prompts the specific constructive processes that *do* take place.To say that perception must be supported by the stimulus is just a different way of saying that the perceptual solution must be 'contained in' the stimulus. But this is precisely the notion we wanted an account of. Hence, the same problem arises: how can perception be contained in the sensory registration, given that the stimulus underdetermines the percept formed on its basis?

The same sorts of problems arise with trying to make sense of the idea that the solution must be 'consistent' with the perceptual solution. What does 'consistency' mean in the context of the relationship between the stimulus and a percept? We can't mean *logical* consistency, since it is *logically consistent* (but not perceptually consistent) with the stick's appearing bent in water that it is in fact straight. I also think we should reject the idea of consistency as accordance with some set of conventional rules for interpreting the image (Wittgenstein uses the metaphor of a 'blueprint' expresses this idea: "For when should I call it a mere case of knowing, not seeing?—Perhaps when someone treats the picture as a working drawing, reads it like a blueprint.") Rather, the relevant notion of consistency seems to be connected to the idea of a g*eometric* '*projection*': the physical stimulus has a geometric structure; and the intentional content of the percept must 'conform' to that structure, in the sense that the permissible perceptions are all and only those ways of 'filling in' the two-dimensional image which are consistent with the laws governing three-dimensional, visual space. The 'building' that perception does is directed by the sensory array in a way that is not conventional but has to do with the geometry of the stimulus. But

to say that perception must 'conform to' or 'respect' the possibilities imposed by the geometry of the stimulus seems again to be another way of saying that the perceptual solution must *be 'contained in' the stimulus*. We are pushed back onto our old question in a new form: how can the percept be 'in' the stimulus, given that the percept is supposed to be something separate from the stimulus? The cognitivist tries to sidestep the problem by discounting the receptivity intuition. Yet on their view, the intuition (and hence the problem) is only pushed back a level.

A rather different source of skepticism concerning the problem I have been describing is the idea that the receptivity intuition is untouched by the problem of underdetermination, despite what I have said. There is a distinction to be drawn between what is attributed to an individual and what is attributed to the individual's subsystems; between the 'personal' and 'subpersonal' levels of description of an individual's psychology. Perceptual experiences and judgements are personal-level phenomena; they are states attributable to subjects, which explain their actions at that level of description. Inferential processes, whether conscious or unconscious, are also personal level phenomena. But the processes which explain the formations of perception from sensation belong to the domain of 'subpersonal processing,' since they are in-principle inaccessible to the subject. Those who champion the Receptivity Intuition most fiercely—naive realists, disjunctivists, McDowellians—have disputed that considerations about the psychological processes by which perceptions are formed are relevant to a theory of perceptual representation at all. They may therefore deny the force of the problem I'm trying to motivate: the considerations which prompt the question pertain to the 'machinery' going on at the sub-

personal level, and so do not bear on the receptivity intuition which pertains to our ordinary, person-level concept of perception. But this is not a satisfactory response. The receptivity intuition holds at the level of the objective 'content' of perception: it concerns the *source* of perceptual objectivity. We have seen that the naive realist intuition about perceptual phenomenology are grounded in this more basic notion of receptivity: how the world seems to us in perception reflects our underlying conception of how perception puts us in touch with the world. Suppose that that underlying conception turns out to be incorrect (as the problem of underdetermination pushes us to think). Then that seeming will be revealed to be in some sense illusory: the directness of perception is *only* a feeling we have about perception; which does not correspond to anything real in the underlying process of generating information. We will have admitted defeat with regards to the true substance of the receptivity intuition. (This issue will be revisited in greater depth in section 4.5.)

To summarize: *the* 'problem of underdetermination' is that, since there is nothing in the structure of the immediate stimulus (for instance, the retinal image) which is specific to its source, we can't explain perceptions purely in terms of their sensory inputs. And this would seem to present a problem for the receptivity intuition, which suggests that *perception can be a receiving of the world only if the outputs of perceptual process are already contained in the sensory inputs, so that we require no internal processing or mental contribution on the part of the subject to explain the formation of perceptions.* We cannot simply sidestep this problem—either by discounting receptivity or by discounting the underdetermination problem. We must find some way to make these two ideas consistent.

243

4.4 Redefining the Stimulus

The first approach we shall consider seeks preserve receptivity in a way that is faithful to our naive conception by revising our conception of the inputs to perception. Hidden in the classical presentation of the underdetermination problem is an assumption about what the immediate stimulus for perception is: the isolated registrations of ambient energy that directly correspond to aspects of our overall perceptual experience. Some philosophers reject the formulation of the underdetermination problem. Merleau-Ponty criticizes the assumption, which he finds in Frege and in cognitivist theories of perception, that perceptual experiences are composites of 'sensations' and abstract propositional content which are constructed from sensations on the grounds that the unit of sensation, 'an isolated datum of perception,' is 'inconceivable' (1945: 4).   Drawing inspiration from Gestalt psychology, he holds that he holds that the subjective and objective aspects of perception are holistically interrelated and that meaning is generated through 'authentic introspection'; "an act which creates at a stroke, along with the cluster of data, the meaning which unites them—indeed which not only discovers the meaning which they have, but moreover causes them to have a meaning" (ibid., 42). By far the most thorough attempts to work out a robust empirical account of this form, J.J. Gibson, is another prominent example of someone who defends receptivity through this approach. Gibson rejects constructivist views of perception precisely for their failure to capture that intuition; for their assumption that meaning comes from 'inside' rather than 'outside,' that objective significance

('meaning') in perception comes from the mind *rather* than the world. "(T)hese theories all assume without question," he writes,

> that sense impressions are somehow the cause of perception but not a sufficient cause. They are taken to be the occasion for perception, the basis for it, or the raw material from which perception is constructed. These theories all take for granted the poverty of the senses and seek for a special process in the mind or the brain to supplant them. They assume that the organs of sense are passive, or merely receptive, accepting whatever physical stimulus enters as if they were merely windows... These assumptions can be challenged ... It might be that no special process is necessary to explain perception, and that in fact perception is not based on sensation (1967: 64).

When Gibson criticizes this 'passive' or 'merely receptive' view of the senses, we should not mistake this for a criticism of the receptivity intuition. Instead, we should take him to be criticizing a view of the stimulus that makes the constructivist view seem inevitable. He is focussing on a different notion of 'receptivity,' one that is contrasted not with 'construction' but with 'disembodiment' or 'disengagement.' He is rejecting a theoretical approach to the senses that ignores the fundamentally embodied character of perception: "The channels for stimulus information that we have arbitrarily separated and called "sensation : are normally active and exploratory, not passive and receptive... (T)he passive arousal of sensations as these have been studied by sensory physiologists is not typically the way

perception works in life" (ibid., 64). Within psychology, 'direct' or 'stimulus theories' of perception which attempt to explain the formation of perceptions from the stimulus considered globally, are as old as inferential views: in the early 20th century, Hering proposed that lightness constancy could be accommodated by appealing to ratios of retinal registrations of luminance. Gibson's 'ecological optics' builds on this by incorporating reference to the ecology of perception. Size constancy / texture gradients. Many surfaces in the world—grassy meadows, brick walls, textured ceilings, and tiled floors—have approximately uniform textures that project texture gradients onto the retina when viewed at a slant. When a surface consisting of many similar objects is slanted in depth, they form an image texture whose elements gradually gets smaller, denser, and more foreshortened as the surface recedes in distance. If we consider only a single brick in this wall, it's retinal size underdetermines its perceived size. But if we consider the brick in the context of the whole retinal image—the texture gradient—this is no longer the case. Emphasizes movement—The same optical information that specifies the nature of the environment also specifies the trajectory of the observer through it. Consider figure 6, below.
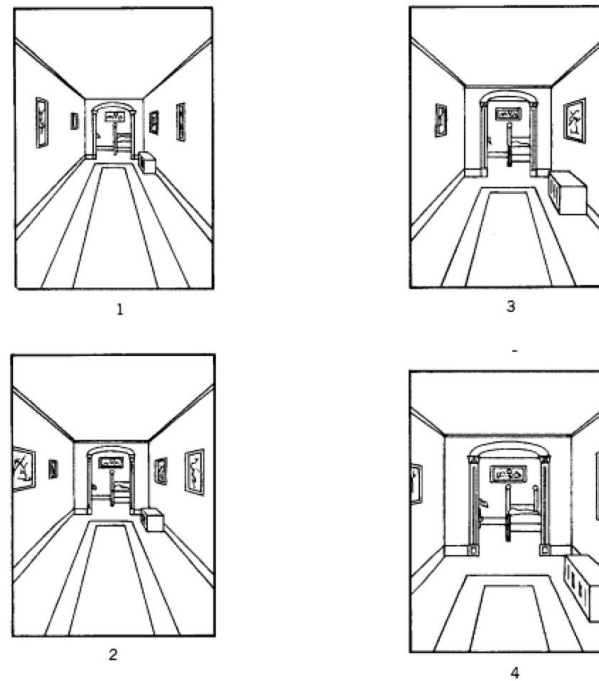
Figure 6. A sequence of four images showing how a figure moving through space receives information that reduces indeterminacy in the retinal image. In this example, the changing size of the doorframe as the subject approaches indicates the effect of distance on retinal size.

The sequence of views shown here not only indicates the presence of a doorway leading to a different room, but also specifies the observer's direction of approach. Perception of the environment and of the moving station point are thus co-determined by the nature of the optical information reaching the eye. Information about the external environment is fully specified by the optical information available at the retina of a moving, actively exploring organism without any mediating processes or internal representations.

Whether this view succeeds in saving perceptual receptivity depends on whether it is correct to claim that the senses are sufficient to account for objective perception. In the light of the empirical evidence, this claim is not viable. Some parts of vision may work this

way—Hering's proposal about luminance ratios is widely accepted, for example—but there's no way to accomodate *all* of the phenomena in 'direct' terms. Perception under less-than-ideal conditions poses a problem. As Roger Shepard points out, those who follow the ecological approach have focussed on explaining "the identification and specification of the invariants that are sufficient for the veridical perception of the local environment under favorable conditions of visibility, mobility, and neural integrity"; yet the stimulus theory struggles to explain perception under "less favorable conditions of nighttime, obstructed, and spatially or temporally limited viewing and, even, of structural damage to the brain itself" (1982, 419). Gibson's view doesn't get the balance of constraint vs freedom right: it errs too much on the side of constraint; and can't succeed in accommodating the full range of the generality and flexibility we find in perceptual capacities to resolve underdetermination in the sensory stimulus. The proposal that perception is a direct pick up of information also struggles to explain the phenomenon of percept-percept causation (Rock, 1997; Palmer, 1997). It is well-established that in certain cases, the same stimulus information can lead to different perceptions of one property, B, when nothing changes but the observer's perception of another property, A. Consider figure 7, below.
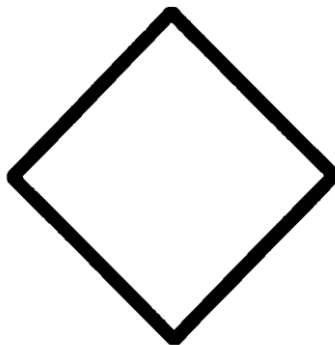
Figure 7. A regular rectangle oriented at an angle of 90 degrees is normally seen as a diamond.

Observers generally perceive this shape to be that of an upright diamond. But if they are told to imagine it as having been tilted clockwise by 45 degrees, so that the upper right segment is its top, they then perceive its shape to be square. Notice that the stimulus information is exactly the same for both perceptions: the only difference is that they perceive its orientation as tilted. But this would seem to be a problem for any stimulus theory because differences in perceptions are supposed to be caused by differences in retinal infonnation: when the retinal information is the same, the perception should be the same. In short, whilst the approach we have considered in this section tries to preserve the naive idea of what it means for perception to be a receiving of the world by overhauling certain assumptions about the sensory inputs to the perceptual process, "the prospects for dispensing with representational content in perceptual psychology are remote" (Burge, 2010: 100fn). These considerations are enough to warrant a decisive rejection of the theory, whatever its other merits may be. In accounting for the problem of underdetermination, there is no getting away from the need for abstraction and internal processing.

4.5 Redefining Receptivity: A Conceptualist Reading of Kant

Although we cannot save the naive concept of receptivity, we may yet save the intuition behind it —that perception comes from the world rather than from the meaning-generating activities of the subject—by redefining what it would mean for perception to be a receptivity. The assumption has been that any kind of information-processing system in

which the outputs are not already contained in the inputs requires some kind of constructive or generative process on the part of the subject to explain how we get from one to the other. Perception can be a passive receiving only if perceptual information is already contained in the sensory inputs, so that no internal processing or mental contribution on the part of the subject is needed to explain the formation of the perceptual states. The foregoing discussion has showed that if there is a satisfying distinction to be drawn between what is 'given' by the senses and what is 'constructed' by the subject, this has to be drawn within a domain of internal processing and 'interpretation.' Is there a way of capturing the core intuition without this framework assumption about what it would *mean* for perception to be a receptivity? This section will consider McDowell's 'Kantian' account of the relationship between perception and thought, argued for in Mind and World, as one proposal for what such a 'reconceived receptivity' might be.

We may begin by noticing that the problem we have been considering is mirrored by a seeming tension within Kant's own account of the distinction between spontaneity and receptivity. The first major section of the Critique, the Transcendental Aesthetic, which is devoted to an examination of the content and origin of sensory representations, begins with the following characterization:

In whatever way and by whatever means a cognition may refer to objects, still intuition is that by which a cognition refers to objects directly, and at which all thought aims as a means. Intuition, however, takes place only insofar as the object is given to us; but that, in tum, is possible only-for us human beings, at any rate-by the

mind's being affected in a certain manner. The capacity (a receptivity) to acquire presentations as a result of the way in which we are affected by objects is called sensibility. Hence by means of sensibility objects are given to us, and it alone supplies us with intuitions. (A19/B33)

The claim is that intuitions (perceptual experiences) are the outputs of a certain kind of cognitive capacity: one whose inputs are always sensations ("the way in which we are affected by objects") and whose function is fundamentally passive ("a capacity to receive presentations") rather than constructive, as thought is ("a capacity to produce presentations ourselves"). Many interpreters of Kant, including McDowell, have understood receptivity in a way that broadly maps onto the naive concept of receptivity: as a purely passive process in which the world 'impresses' or 'impinges' itself onto a possessor of sensory capacities. The concept of intuition has its roots in medieval philosophy, where it was used to describe knowledge acquired directly, without needing to be worked up through some process. Kant is essentially affirming the existence of this kind of knowledge, whilst restricting it to the domain of perception (hence denying, for example, the idea that we receive our knowledge of moral truths in this way, as someone like Augustine claimed). It is specifically objective perceptions—rather than mere sensations—which are being produced through the 'receptive' process. Kant's claim in this passage thus reflects our naive intuition that perception is a distinctive kind of capacity for acquiring information about the world in virtue of the fact that that information is simply 'given' by the effect objects have on us, without our having to do anything to bring them about.

251

The puzzle is that this would appear to be in tension with certain commitments Kant goes on to express immediately following this passage, when he discusses the distinction between sensation and intuition:

> The effect of an object on our capacity for presentation, insofar as we are affected by the object, is sensation. Intuition that refers to the object through sensation is called empirical intuition. The undetermined object of an empirical intuition is called appearance. Whatever in an appearance corresponds to sensation I call its matter; but whatever in an appearance brings about the fact that the manifold of the appearance can be ordered in certain relations I call the form of appearance. Now, that in which alone sensations can be ordered and put into a certain form cannot itself be sensation again. Therefore, although the matter of all appearance is given to us only a posteriori, the form of all appearance must altogether lie ready for the sensations a priori in the mind; and hence that form must be capable of being examined apart from all sensation. (B34/A20)

The claim that some features of our knowledge are a priori means, in this context, that they do not derive from sensory evidence, but from our minds' ways of dealing with sensory evidence (as he presents this idea in the introduction to the Critique, "In part these objects by themselves bring about presentations. In part they set in motion our understandings activity, by which it compares these presentations, connects or separates them, and thus processes the raw material of sense impressions into a cognition of objects that is called

experience.") It is clear from this passage that Kant regards perceptions—sensory representations having objective significance—as requiring a contribution from the mind. Intuitions are distinguished from mere sensation by virtue of containing an a priori element which is *not* derived from the senses: something which "lies ready for the sensations … in the mind"; and by which sensory matter is "ordered and put into a certain form". Some have interpreted this as speaking to Kant's recognition of the problem of underdetermination: Kant's point is that "the perceptual faculty must interpret the two-dimensional information on the retina as a three-dimensional array… the sensory information registered on our visual organ constrains, but does not determine, a three-dimensional visual image; the production of that image requires work from the faculty of perception. In contemporary terminology, Kant's claim would be that the sensory data on our retinas must be processed by our perceptual systems before the visual perception of a house, for example, is possible" (Kitcher, 1996: xxxiv). The question is how Kant can consistently maintain these claims in the light of his commitments concerning receptivity. How can it be true that intuitions are both produced from sensation by a fundamentally passive operation, yet also have a structure which involves a mental contribution which organizes sensory intake into a novel product. How can perception be both a passive receiving and an organizing at the same time? In addition to this puzzle, there is the further point emphasized by McDowell that such a view would seem to be in tension with the epistemic concerns which motivate Kant's investigations in the Critique. One of the main points of giving a role to receptivity is to explain how our thoughts connect up with the world ("for I can think whatever I please, provided only that I do not contradict myself").

But talk of a subject's experience as "made up of impressions, impingements by the world on a possessor of sensory capacities" offers us 'exculpations where we wanted justifications.' For justifications can hold only within the space of reasons; whereas "talk of impingements by the world is 'empirical description'... the idea of receiving an impression is the idea of a transaction in nature." If perception is a 'brute' given—something without representational content, hence outside of the 'space of reasons'—it would be useless for the purpose for which Kant introduces it. The idea that Kant may hold a solution to our problem seems natural when we consider this tension. Whilst it is of course possible that Kant was simply unaware of these tensions, or had no answer for them, the more likely scenario is that he is thinking of receptivity in some way that is different from the naive conception.

McDowell's proposal for how to diffuse this apparent tension lies in Kant's views about how spontaneity and receptivity interact in the formation of perceptions. The key source for this interpretation is a discussion in the Deduction in which Kant emphasizes that empirical knowledge results from a mutual interdependence of receptivity and spontaneity: "Thoughts without content are empty; intuitions without concepts are blind. Hence it is just as necessary that we make our concepts sensible (i.e., that we add the object to them in intuition) as it is necessary that we make our intuitions understandable (i.e., that we bring them under concepts). (B76/A52). On the perception side of this equation, this amounts to the claim that experiences result from a cooperation between receptivity and spontaneity —where "Receptivity does not make an even notionally separable contribution to the co-operation." That is,

The relevant conceptual capacities are drawn on *in* receptivity… It is not that they are exercised *on* an extra-conceptual deliverance of receptivity. We should understand what Kant calls "intuition"—experiential intake—not as a bare getting of an extra-conceptual Given, but as a kind of occurrence or state that already has conceptual content… In the view I am urging, the conceptual contents that sit closes to the impact of external reality on one's sensibility are not already, *qua* conceptual, some distance away from that impact. They are not the results of a first step within the space of reasons, a step that would be re-traced by the last step in laying out justifications, as that activity is conceived within the dualism of scheme and Given. This supposed first step would be a move from an impression, conceived as the bare reception of a bit of the Given, to a judgement justified by the impression. But it is not like that: the conceptual contents that are most basic in this sense are already possessed by impressions themselves, impingements by the world on our sensibility. (1994: 24)

For our purposes, the key feature of this passage is its abandonment of the naive view of receptivity, in favor of a reconception of what it means for perception to be a receptivity. The abandonment of the naive view is marked by its rejection of "a bare reception of an extra-conceptual Given." Perception is a receptivity, but it is not *just* a receptivity: the operations of receptivity by themselves are not sufficient to generate objective perception. Perceptual experience is partly a product of spontaneity; which in the present context

means that perception involves conceptual capacities: the contents of perception are the very same things that form the direct objects of judgements and other propositional attitudes.[43] The proffered reconception of receptivity is the other side of this coin: if receptivity cannot operate in isolation from spontaneity, then objective perceptions cannot be a construction out of 'bare impressions.' Whereas judgements are produced out of experiences, perception is not a construction out of something already given. We should not think of sensory registrations as an initial datum for subjects from which they make a step to a perceptual experience; sensory matter exists only in unison with an objective (conceptual) content. This is not to deny that there is a separation between the immediate stimuli that are registered by the senses and the experience that is ultimately generated from those registrations. But this process will be at a 'subpersonal' level, and not something performed by the subject. As Tyler Burge has put it, "The system transforms registrations into perceptions through a series of stages. The animal does no such thing. So the perceiver's perceptual representations— though they are also products of the perceptual system— are not the result of the perceiver's activities, only the results of his perceptual system's transformations. In this weak sense, the perceiver's representations are (relative to the perceiver's activities) transformationally null ("direct"?), whereas (relative to the

---

[43] Notice that this tendency to downplay the role of receptivity—viewing it as capable of bringing about perception only in conjunction with operations of spontaneity —is not unique to conceptualist interpreters of Kant. Robert Hanna, who favors a nonconceptualist view of Kant, takes the same view: "To regard sensibility as wholly passive would be mistakenly to identify it with Locke's model of sense perception—the mind as a sort of conscious black box with a peephole to let in the light, and an impressionable blank tablet on the inside; a mental camera obscura. By sharp contrast, for Kant the sensibility and the understanding alike are generative and productive sources: each is a cognitive 'capacity' … for spontaneously simplifying and interpreting—for spontaneously informing and transforming—inputs" (2001: 37). Hanna accommodates this point within a nonconceptualist framework by drawing a distinction between conceptual and nonconceptual forms of spontaneity: "the most accurate way of characterizing the relationship between sensibility and understanding is in terms of two distinct levels of spontaneity of synthesis: lower-level or sensory (receptive) spontaneity; and higher-level or conceptual (discursive) spontaneity."

perceptual system's transactions) they are transformationally complex ("indirect"?)" (2005: 49).

According to McDowell, this qualified role for receptivity gives us "all the external constraint we need." That is,

> The fact that experience is passive, a matter of receptivity in operation, should assure us that we have all the external constraint we can reasonably want. The constraint comes from outside thinking, but not from outside what is thinkable. When we trace justifications back, the last thing we come to is still a thinkable content; not something more ultimate than that, a bare pointing to a bit of the given. But these final thinkable contents are put into place in operations of receptivity, and that means that when we appeal to them we register the required constraint on thinking from a reality external to it. (1994: 28-29)

For Mcdowell, the necessary role for spontaneity pertains to the need to recognize contentful—and not merely causal—relations between experience and thought in a sound epistemology. But one can also see how it might be extended to the kinds of issues with which we have been most centrally concerned. In insisting on a role for conceptual experience, … Allows for subpersonal processing—allows us to resolve … But "perceptible facts are essentially capable of impressing themselves on perceivers … and facts in general are essentially capable of being embraced in thought in exercises of spontaneity" (1994: 28). This proposal also seems to give an explanation for the intimate relationship between

the subjective and objective components of experience. Dan Crawford—who agrees with McDowell not only in the view but also the claim that it is Kant's—thinks the main objection to theories of perception that regard perception as a construction (or inference) out of sensory registrations is that our sense impressions are not always on the surface of, and discoverable in, conscious experience. We must recognize three separate causal stages of the perceptual process: the purely physical stimulus, which gives rise to the mental sense impression, which in turn gives rise to the "conceptually rich" perception. But "it will not do to say simply that sense impressions causally mediate perceptual awarenesses, for they have a far more intimate connection with those awarenesses. We can only account for the sensuous, non-propositional character of perceptual states, if we accept a strong form of the Kantian doctrine that the sensory "matter" is literally taken up into the perceptual experience and seamlessly united with it although … it can be significantly altered in the process" (1982: 88).

Nevertheless, I do not think that what McDowell is proposing is a truly satisfying solution to our problem. In a way, it captures the surface of the intuition but not the core. The worry here goes back to the problem I discussed in section 4.3, when I was motivating the problem of perceptual receptivity. For McDowell, the details of the process that generates perceptual experiences out of sensory registrations are not important. That work is at the 'subpersonal' level whereas receptivity is a personal-level phenomenon, consisting in an experience in which conceptual content and sensory matter are seamlessly combined. But whilst this explains why perception *seems* to us to be immediate and direct, receptivity is vindicated only at a superficial level. The synthesis is achieved *post hoc*: it is the sensory

matter that is taken up into the conceptual content and altered by it; rather than the content coming out of a sensory matter that contains it. McDowell claims that this is "all the constraint we need." Yet at an underlying level, the direction of explanation for our acquisition is no more world-to-mind than it is for the case of inference. The propositional content is still fundamentally separate.

I also think there are reasons to be dissatisfied with this as a solution to Kant's problem. Consider the following:

> Only from the human standpoint, therefore, can we speak of space, of extended beings, etc. If we depart from the subjective condition under which alone we can-viz, as far as we may be affected by objects-acquire outer intuition, then the presentation of space means nothing whatsoever. This predicate is ascribed to things only insofar as they appear to us, i.e., only insofar as they are objects of sensibility. The constant form of this receptivity which we call sensibility is a necessary condition of all relations in which objects are intuited as outside us; and if we abstract from these objects, then the form of that receptivity is a pure intuition that bears the name of space. (A26-27/B42-43)

Kant's avowals (here and at B44) that space is 'the form of receptivity' seems fundamentally at odds with McDowell's conception of the interplay between receptivity and spontaneity. For McDowell, receptivity is a 'bare getting of the given.' Within McDowell's framework, the idea Kant is expressing here would be understood in terms of

the idea that the sensory and conceptual elements of perception are inseparable from one another: the form of receptivity that is a necessary condition for objective experience is a conceptual content of some sort. But this seems straightforwardly incompatible with Kant's view that the form of receptivity makes no sense independently of the subjective aspect: "We cannot make the special conditions of sensibility to be conditions of the possibility of things, but only of the possibility of their appearances. Hence we can indeed say that space encompasses all thing that appear to us externally, but not that it encompasses all things in themselves, intuited or not, or intuited by whatever subject." The interdependence of the two elements seems to be greater than McDowell acknowledges. Despite these concerns, I think McDowell is right in suggesting that Kant can help to resolve the problem of perceptual receptivity by giving us a different way of thinking about what it is to 'receive' presentations of the world. In the next section, I shall articulate a rather different view about that is.

4.6 Redefining Receptivity: A Nonconceptualist Reading of Kant

Kant holds that both sense perception and conceptual thought contain a priori elements. Alongside his distinction between two faculties, Kant identifies two kinds of a priori cognitions corresponding to these: one, which we have already met, are the 'pure forms of sensibility'; the other are the 'pure concepts of the understanding' ('pure,' here, is Kant's term for representations containing no 'empirical' element mixed in with them). I have noted the temptation that exists to move immediately from this to the idea that perception, like conceptual thought, involves a mental *processing*; an activity or work that the mind

performs *on* something. Yet I believed this is to overlook what is amongst the most creative

aspects of Kant's account of sensory representation, namely that these two notions—that of

organizing or giving coherence to sensory registrations on the one hand, and that of

'processing' or 'producing' on the other—can come apart. Kant's concern with identifying a

priori elements in cognition, it is well known, is a concern with identifying what must

already be in place within us for knowledge about the objective world to be possible. But

we can be more specific than this. I believe that the role of the a priori, in this context, has

to do with accounting for the source of the *questions* to which processes of knowledge

acquisition are an answer. To explain the operation of both capacities in generating

representations of the objective world based on subjective sensation, there must already be

in place something non- experiential which imposes 'principles of interpretation.' In

distinguishing a priori *intuition* from a priori *concepts*, however, Kant is drawing our

attention to the fact that *the principles that are in place* are in each case quite different.

What makes perception a receptivity is not that it does not contribute anything to what is

taken in through perception, but that it does not involve *doing something* to sensory

registrations. Perceptual capacities are set up in such a way that the organizing structures

proceed the inputs: perception does not operate on sensory impressions to generate

perceptions; rather, the receiving of impressions is also, at the same time, a receiving of

perceptions.

Let me begin by elaborating on this idea with reference to the 'pure concepts of the

understanding,' a list of twelve concepts that lie a priori in the mind which Kant calls the

'categories' (A80/B106). The categories are presented as forming a single exhaustive list,

with the four classes of categories imposing four different forms of unity on the object

known: (1) Quantity (Unity, Plurality, Totality); (2) Quality (Reality, Negation, Limitation);

(3) Relation (Inherence and Subsistence or substance and accident, Causality and

Dependence or cause and effect, Community or reciprocity); and (4) Modality (Possibility,

Existence, Necessity). The four major subtypes—quantity, quality, relation, and modality—

represent four independent dimensions by which one can investigate an object's nature,

receiving one of the three sub-answers in each case on the way to a more complete

characterization of the object. Kant's categories closely follow Aristotle's system of

categories—except that where Aristotle aimed to provide a complete list of highest kinds

or genera that exist, Kant's project is to provide an inventory of everything there is *for us*.

That is, much as P. F. Strawson's "descriptive metaphysics" aimed to describe "the most

general features of our conceptual scheme" (1959), Kant aims to uncover the source of

these metaphysical kinds in principles of human understanding; the innate system for

organizing existence that subjects must have to explain the process of acquiring knowledge

in thought. To illustrate, suppose you see the lights on in your neighbor's house and

conclude that the neighbors are home. How are we to explain your acquisition of this

(correct) information? Part of the explanation lies in your capacities to draw the

appropriate inferences from the perceptually given information: you recognize that what

you are seeing are lights are on in the neighbor's house; you generate various hypotheses

about what might have caused this; and you select from amongst these the most likely one.

But whilst much of this process can be explained in terms of the conceptual knowledge we

build up through experience—our acquisition of concepts by which we classify things are

abstracted from perceptual experience and so is our knowledge of how categories are related to one another—Kant's point is that the concept of a cause (which we draw on in formulating the question which we then go on to answer) could not have been derived from experience. Kant agreed with Hume that the concept of cause cannot be traced back to particular features of sensory experience—all you ever see is one thing followed by another—although he disagreed with the conclusion Hume drew from this that such a concept could not legitimately be applied to the data of the senses. His goal in the Transcendental Deduction was to show the legitimacy of this and other non-empirical concepts by showing that they are indispensable for any cognition at all. For example, the claim that all events have causes is universally and necessarily true of all the events of which we can have any cognition (A92-93/B124-126).

I am less interested in assessing Kant's specific claims about the a priori nature of the categories than I am in understanding their implications for our understanding of the structure of problem solving in thought. I believe Kant's distinction between a priori and empirical elements of thought draws our attention to a salient and centrally important distinction in understanding the problem-solving process in thought. One aspect of the problem-solving process has to do with all the knowledge that we accumulate about the structure of the world: the existence of things like burglars and lights and neighbors; and how these things are related to one another. This aspect of thinking is one thing that makes thinkers smart and some thinkers smarter than others (adults are better at figuring out what the world is like than children, and some adults are better than others, because they have more knowledge about the structure of the world). But underlying all this activity is a

separate and more fundamental ability: that of generating the problems or questions that these activities are designed to solve. Before you can search, you need to have some idea of *what you are searching for*. To have generated a classification of the perceived object—with its constellation of colors, shapes and so on—as a cell phone, I must first have thought to ask 'what kind of thing is that?' To have generated, as an explanation for its presence on the table, the hypothesis that my partner came home and left it there, I must first have thought to ask, 'what caused this to be here?' Although this capacity to generate questions plays a role that is less obvious than the capacities we draw on to solve those questions, it is an equally indispensable—and no less remarkable—aspect of the intelligence of thought. What distinguishes the good detective from the incompetent one, the extraordinary scientist from the mediocre one, is not just the breadth of knowledge about how the world works which they bring to their problems, but their ability to recognize where a problem exists; where there is a question to be answered. (Of course, the two are intimately related: to know that there is a problem involves recognizing that things do not quite add up.) Part of Kant's point, in introducing the categories, is to say something about the source of this second kind of ability. What 'lies at the basis' are a list of very general metaphysical categories. Inasmuch as these are a priori, they are fixed: they cannot be added to, or subtracted from, save by some kind of evolutionary leap which changes the internal constitution of our minds. What I take from Kant's idea—whether or not he is right about the categories being innate—is that our ability to discover what the world is like is constrained by our understanding, at a very general level, of the kinds of things that exist. If there were another category, not in Kant's list, we would not be able to discover facts about

entities in that category, even if it were right under our noses, because we would not know to ask about it. The limits of our capacity to acquire knowledge are set by our knowledge of what to look for.

Let us now turn to the forms of sensibility; in particular, the form of 'outer sense,' which Kant identifies as the representation of 'space.'

> Space is nothing but the mere form of all appearances of outer senses; i.e., it is the subjective condition of sensibility under which alone outer intuition is possible for us. Now, the subject's receptivity for being affected by objects precedes necessarily all intuitions of these objects. Thus we can understand how the form of all appearances can be given in the mind prior to all actual perceptions, and hence given a priori; and we can understand how this form, as a pure intuition in which all objects must be determined, can contain, prior to all experience, principles for the relations among these objects. (B43/A27)

We may distinguish two ways of taking Kant's claim that space is a subjective condition. On the one hand, Kant is making a *metaphysical* claim about the nature of space: space is 'transcendentally ideal' in that it is a property of subjects rather than being a property of mind-independent reality; it is "nothing" as soon as we omit reference to the subject and we are wrong to suppose space to be something underlying things in themselves. But within this is a psychological claim about the nature of our perceptions of space and spatial properties: spatial perceptions are *appearances* rather than the thing-in-itself; and "if we

annul ourselves as subject, or even annul only the subjective character of the senses generally, then this entire character of objects and all their relations in space and time-indeed, even space and time themselves would vanish; being appearances, they cannot exist in themselves, but can exist only in us." We do not have to accept the metaphysical claim in order to draw insight from the psychological claim that space a necessary element of intuition. Like the categories, space dictates at a most general level the kind of information we can look for in seeking to characterize an object. It lays down a general system of principles, for perception, concerning what there is: three dimensions of space; the values an object may have along each of these dimensions; and the higher order properties pertaining to the relations of objects to that space (the dimensions of size, shape, orientation and so on). It is fixed, and things that come in are placed *within it*. Kant's views reflect an idea Roger Shephard expresses in in terms of perceptual systems 'internalizing' the 'set of constraints that govern three-dimensional Euclidean space: "What is internalized at the deepest and most abstract level is not any particular object or transformation (which are arbitrary with respect to orientation and path) but the set of constraints that in three-dimensional Euclidean space govern the possible projections and transformations of an object" (1984: 442). The a priori form of space determines what kinds of things our perceptual systems are capable of 'looking for' in the sensory image: it lays out in advance that any sensed item will have a size, a shape, a location, an orientation, and so on. If the existence of these 'kinds' were not already present in the mind, we could never have known to look for them at all.

As a priori systems for organizing perceptual data, however, the categories and the forms are importantly different. Whereas the categories exhaustively dictate the kinds of things there are at a general level, they do not dictate the specific properties we can discover and assign to objects: what is not fixed by the categories are all the kinds that fall under the headings of kinds of substances, kinds of causes and causal laws, and so forth. Compare the case of space, which is 'essentially one'; a *system of magnitudes* with a fixed number of dimensions, and a fixed range of properties which result from higher-order properties of an object's relation to space. Inasmuch as space lies at the basis, all the kinds of properties are fixed in advance. What is open to discovery is only an object's specific values along each of these dimensions. We could not discover that a visually perceived object did not have a third dimension; that it had a width but no height; that it had a shape but no orientation in space. The second dimension of the problem-solving process in thought that I identified above is missing in perception. The world of space is this one single thing: all of the dimensions are fixed. So you can't add a dimension or property: you can't discover a fourth dimension to visual spatial experience; or add a kind of visual spatial property alongside shape, orientation and so on. Nor can you take those things away. And the reason for this is that those properties are just aspects of the system, and the system is essentially whole, 'essentially one.' A further difference between them is that space not only dictates the properties about which perception can 'enquire'; it also dictates the laws governing these kinds—which are the laws of euclidean geometry which govern that system. In the case of thought, you can change your understanding of laws: you can add to them or modify existing ones. You can learn novel laws about what might cause the

lights to be on in a neighbor's house; or what might cause a mobile phone to be on the table; or about the conditions under which you can take what someone says at face value (and the conditions under which someone is likely to be lying). But a spatial system dictates not only the variables but the ways in which all the laws governing those variables are related. For example, the laws relating retinal size, distance and size are fixed by the system, and could not be changed, or updated, or added to. Moreover, it is not just the kinds that are fixed in advance, but also all the laws relating kinds. In thought, we apply our knowledge of rules—of the structure of the world—*to* an object that is given in perception. We can choose to apply this or that rule, or change our understanding of what the laws are.

I think what Kant meant by calling perception a 'receptivity' pertains to these differences between the categories and the form of space as sources of organizing principles. Perception, like thought, generates objective information through a system of internalized principles about what there is and how these things are related to one another. But in perception there is nothing *to which* this system is applied; it is just *there*, as the 'constant form of our receptivity.' That this form of space 'precedes' the sensory inputs that are placed 'within' it does not mean that there is no work for the perceptual system to do in order to figure out where those sensory inputs belong in such a system. Yet that work is all in the service of the predetermined questions that are fixed by the system. *All the questions to be answered have been asked in advance*. And where inference explains the freedom and flexibility we find in thought, receptivity—this mediating system of spatial properties and rules which is fixed and 'constant' and in which all incoming sensory stimuli are assumed to

have a place—predicts just the kind of limited flexibility that we in fact observe in perception.

4.7 The 'Predictive Processing' Theory

The idea we got from Kant is that perception is a 'receptivity' in the following sense: even though there is *processing* of the sensory stimulus, there's a sense in which the processing precedes the inputs. Thought starts from scratch: you get your perceptual data and you apply a strategy or process of knowledge-acquisition. There is a freedom to this: you can decide to apply this approach or that one. But in perception, those procedures are somehow fixed in advance. Moreover, they are fixed by the 'forms' of space and time. These are promising ideas, yet they remain largely suggestive. What's needed is a theory of perception which illustrates in a concrete way how Kant's claims are realized in the resolution of the underdetermination problem. This section will argue that 'predictive processing' theories of perception, developed in recent decades by (Hinton, 1983; Jaynes, 1988; Dayan, 1992; Watkins, 1992) may provide just such an illustration.

To get a sense for the theoretical framework behind predictive processing models of perception, begin by considering an example of this strategy in action in belief. Imagine a subject who tests positive for a disease. The test is 99% accurate. For many people, their first thought is likely to be that they have a 99% chance of having the disease. But this fails to take into account the rarity of the disease, which occurs in only one in 10,000 people. The test gives a false positive 1% of the time. And for such a rare disease, there would be many false positives in a given population. Thus, your actual likelihood of having this

disease is much lower than you initially supposed. What this example illustrates is a way of optimizing (making more accurate) your subjective beliefs through the addition of statistical information which takes into account the objective probabilities of the various elements of the scenario (the reliability of the test, the likelihood of the disease). Notice that Bayes' theorem is not meant to be a description of how we actually think (although it takes as its initial priors the beliefs we arrive at through inference). In fact, it is quite counter intuitive. It's value lies in the fact that, through an essentially iterative model—one in which priors, together with statistical probabilities, generate new more accurate information which can in turn be used in the same procedure—it is able to achieve more accurate beliefs and more effective decision-making procedures.

The power of this model, put to work in a practical domain, lies in its ability to resolve information processing problems using a simple, iterative procedure which presupposes relatively little in the way of processing capacities (certainly nothing like the machinery of concepts, knowledge of generalizations and rational capabilities presupposed in inference). Consider data compression strategies such as image transmission. Data compression works through an algorithm which makes assumptions about the likely value for one pixel, given its neighboring values (simplest prediction would be that neighbouring pixels all share the same value (the same grey scale value, for example). The only information that needs to be transmitted is that concerning values which are *unexpected*; which deviate from the predicted value (for instance the boundaries between objects). This affords major savings on bandwidth. Data-compression by informed prediction allows quite modest encodings to be reconstructed into rich and florid renditions of the original

sights and sounds. Such techniques figure prominently in, for example, motion-compressed coding for video.

At the core of the Predictive Processing theory of perception is the idea that this framework supplies us with a model with which to understand how perception resolves underdetermination in the sensory stimulus to give us an objective world. The basic form of the information-processing problem for perception, we shall recall, is that of getting from the flows of energetic stimulation that impinge upon sensory apparatus to the mind-independent objects and properties that caused those energetic stimulations. We know that perception *does* achieve this. The question is how. If all the system ever has direct access to is its own sensory states—patterns of registrations of light on the retina, say, or patterns of registrations on the nerves of the surface of the skin, how do we ever get a grip on a world of objects with shapes and sizes, colors, locations and trajectories through space, and infinitely many more properties besides? The answer, on a predictive processing model, is through *prediction*. From the day you are born, your brain begins making predictions about what will happen, based on what *has* happened—and as it makes mistakes, it learns from those mistakes in order to generate new and better predictions. Predictive processing models of perception combine the theoretical model of predictive coding with a further element of learning procedures. Thus:

> predictive forms of learning are particularly compelling because they provide a
> ubiquitous source of learning signals: if you attempt to predict everything that
> happens next, then every single moment is a learning opportunity. This kind of

pervasive learning can for example explain how an infant seems to magically

acquire such a sophisticated understanding of the world, despite their seemingly

inert overt behavior (Elman, Bates, Johnson, Karmiloff-Smith, Parisi, and Plunkett,

1996)

The suggestion here has elements in common with the Ecological approach, in the sense

that it emphasizes the fundamentally embodied nature of the perceiver and the effect of

this on the sensory input. The task is to infer the nature of the signal source (the world)

from just the *varying input signal itself.* Because of this embodied character of experience,

the world itself is able to provide a "training signal." "For the states of your sensory

registers will change, in ways systematically driven by the incoming signal, as the world

around you changes. In this way, the evolving states of your own sensory receptors provide

a training signal allowing your brain to 'self-supervise' its own learning."

It may seem odd to speak of the perceptual process in this way. Prediction is future

oriented. But the more natural way to think about the underdetermination problem is past-

oriented. Think of the case of Sherlock Holmes. It seems bizarre to speak of him as

predicting who committed the murder, based on the perceptually available clues. Rather, he

uses those clues as evidence from which to reason backwards to a picture of what caused

the effects that are perceptually available to him. Likewise, it is tempting to think of the

problem for perception as a matter of retracing or working backwards from causal effect to

causal antecedent: from sensory registrations to the distal states of affairs that gave rise to

them. But this reorientation is precisely one of the key insights of the predictive processing

272

model, showing how perception is able to resolve a problem that 'looks like' the kinds of problems we are able to resolve in thought, through a process that is *completely different* in its underlying form. It is also one of the ways in which the predictive process links up with Kant's idea of perception as a 'receptivity.' Given that certain conditions are met (of which more anon) it is possible to achieve the same result that would be achieved through the conceptualized structure of "cause and effect" through an entirely unintelligent, 'thoughtless' process based in *statistical regularities*. Your first prior will just be whatever has already happened. Once something happens that deviates from this prior, you change the prior: it is capable of changing; and the likelihood of its changing is X. At first it would be very basic: that everything would stay the same. When things change you will update this. You start to look for patterns: regularity, predictability. At the start (or maybe all the way though), these predictions don't need to be at the level of a person's awareness—(this wall is going to stay the same. People are going to move. Shadows are going to move. Surfaces are not going to change color (most of the time). Things are not going to change size. The source of light changes. But notice that these regularities (or 'heuristics') are not encoded as personal level, objective representations as they would be in thought. They are just reflections of statistical data—of the kind that could be encoded by a fairly simple machine. As Andy Clark puts it, the prediction involved in perception "*is the kind of automatically deployed, deeply probabilistic, non-conscious guessing that occurs as part of the complex neural processing routines that underpin and unify perception and action.*" What is remarkable, on this view, is not how 'intelligent' perception is, but how much of the world can be built up from such a simple basis.

Interestingly, when we view predictive processing as a concrete articulation of Kant's theory of receptivity we get a very different view of the relationship between perception and cognition than the view Baysians themselves have often reached. Many Baysians have been skeptical about the idea of a joint in nature between perception and cognition distinction, holding that that within the framework of predictive processing, the lines between perception and cognition become "fuzzy, perhaps even vanishing" (Clark, ibid., 10). I am suggesting that predictive processing models actually help to clarify how there are two fundamentally different ways in which the mind may bring structure and order to what is taken in through the senses. To come back to the idea of a problem-solving process: to solve a problem requires that one see a problem to be solved; and to see a problem "presupposes, by definition, the awareness of an obstacle which does not allow to reach a certain goal…. Problems do not exist in nature; they exist only when there is a mind which experiences a situation as problematic" (Kanisza, 1985: 27). The model of predictive processing shows in a concrete way how we credit perception with too much intelligence when we try to apply the idea of 'problem-solving' in a way that goes beyond a purely metaphorical usage. In order to determine the specific value that a sensory stimulus has along each of these dimensions, the visual system must use 'heuristics' that in some ways echo the generalizations we make use of in thought; yet the process and principles of discovery for perception are 'automatic,' dictated by the accumulation of statistical regularities. There is no problem-solving, because there is no recognition of a problem *to be solved*. Perception can interpret only through arranging sensations within a pre-programmed framework that is fitted to receive them, and that is very different from the

274

deliberate process of searching for explanations by which inferential thought achieves its goals.

4.8 Conclusion

This chapter has argued for two substantive conclusions about the nature of perception. The first is that perception is distinguished from thought by being a 'receptivity to the world'—on a particular understanding of this claim. If we are to bring receptivity into alignment with the fact that sensory registrations underdetermine the objective contents of perception, we cannot maintain a view of receptivity as a purely passive operation, as some have wanted to claim. But nor must we do so, in order to preserve what is truly important in the receptivity intuition. Following Kant's suggestion that the forms of intuition (notably, space) are *forms of receptivity which lie innate in* the mind and proceed the data of the senses, I have argued that perceptual receptivity consists in perception's possessing a certain *kind* of organizational structure; one in which the registration of sensory information is also, and already, a receiving of objective information. Unlike the rational and conceptual schema governing thought, the fixed and systematic organizing principles which are distinctive to perception use sensory registrations in the generation of meaning. The second conclusion, a consequence of and a compliment to the previous one, is that perception has a nonconceptual content: a non-propositional content in which the real of sensation makes a semantic contribution to perceptual content. Underlying these substantive conclusions is a broader aim concerning the status of the receptivity intuition in theorizing about perception. I have tried to show that the receptivity intuition is central

to both commonsense and theoretical thinking about perception. In addition to its explicit mention by philosophers such as Kant, McDowell and Burge, the idea that perception is distinguished from thought by being a 'receptivity to the world' plays an important (though often unacknowledged) role in theories of perception put forward by naive realists and cognitivists. If nothing else, I hope my reader will come away from this discussion with a sense that the concept of perceptual receptivity is worthy of greater attention and more careful elucidation than it has generally been accorded in recent discussions.

In exploring our intuitions about the perception/cognition distinction, I have focussed on certain sorts of examples in which those intuitions manifest themselves: examples in which we make judgements about properties of objects which, under other circumstances, we could have perceived (the yellowness of a lemon or its shape; the blonde of a boy's hair). These kinds of cases are one way of exploring the question of what the criteria for perception are: in virtue of what do we count something as a perception rather than a thought? But they have not been the most common kinds of cases philosophers have considered when addressing this question. Consider the following:

> Is it already part of your visual experience that John Malkovich is walking by, carrying a dog? Or do you just visually experience an array of colored shapes bouncing slightly at regular intervals, and subsequently judge that it is John Malkovich carrying a dog? More generally, we can ask: do you just visually experience arrays of colored shapes, variously illuminated, and sometimes moving? Or does visual experience involve more complex features, such as personal identity, causation, and kinds such as bicycle, keys, and cars? (Seigel, 2006: 3)

Is there a viable distinction to be drawn between perception and cognition? There certainly seems to be a difference in kind between hearing a balloon pop and thinking about the square root of −1. (Phillips, 2017: 1)

These are questions about the *admissible contents of experience*: which kinds of properties are part of our perceptual repertoire and which can we say are definitively cognitive? Is there any principled basis on which we can say this at all? **One attractive** feature of the approach I have adopted in this chapter lies in its promise to supply answers to this second sort of question, by giving us a principled basis for sorting through those cases where our intuitions are less clear-cut. Perceptual receptivity marks a particular kind of process for generating representations; one that might correspond to predictive processing models of information-processing, and which contrasts with inferential models. These two levels of representation—perception and cognition— although they may represent the same properties, involve different psychological competencies and utilize different types of contents. Assessing whether a borderline representation is represented in perception becomes a matter of examining the patterns of success and failure predicted by the underlying mechanisms. Such a proposal compares favorably with other approaches, such as Seigel's 'method of phenomenal contrast,' insofar as it makes the question one amenable to study within a scientific psychological paradigm. I agree with Burge in his view that phenomenological and other armchair approaches to determining which properties are represented in perception are "deeply wrong-headed": "One cannot distinguish cognition from perception in any warranted way from the armchair. The processes for forming attributives on the basis of perception are too fast, inaccessible to consciousness, and

complex to allow phenomenological or other armchair methods to distinguish perception from cognition. Only sophisticated use of experimental evidence bears on these issues in a way that goes beyond uninformed playing" (Burge, 2014: 11).

Although my conclusions primarily address the nature of perception, they also have implications for Kantian scholarship. In claiming Kant as a source of inspiration for my proposals, I am taking a stand in the debate over whether Kant held a nonconceptualist view of perception—a debate that has come to seem almost as intractable as the nonconceptualism debate itself. But a notable feature of the present account is to present this question as turning on issues quite different from those generally considered. I reach a nonconceptualist reading of Kant by challenging a more fundamental assumption which these disagreements have generally presupposed: that Kant's opposition between receptivity and spontaneity is a distinction between a purely passive registering inputs and processing of inputs that involves contributions from the mind. In my view, this 'sidelining' of receptivity overlooks the features of Kant's theory of perception that are the most creative and the most valuable to contemporary philosophers of perception.

The first half of this dissertation was concerned with understanding the positions that are at stake within the nonconceptualism debate: chapter one outlined the contours of the disagreement between the proponent of nonconceptual content, the conceptualist, and the naive realist; whilst chapter two defended nonconceptual content against charges of incoherence by those on either side. This chapter, particularly when taken in conjunction with the preceding one, motivates a resolution of that debate in favor of the proponent of nonconceptual content. Starting from an intuition which is common currency amongst all

three parties to this debate—the idea that perception is a receptivity to the world—I have argued that our best option for making sense of that claim entails a nonconceptual content for perception.

CONCLUSION

If we deny that perception is reducible to thought, should we also deny that perceptual presentations can be captured by the contents we use to characterize thoughts? Does the distinctive character of perceptual states entail a distinctive kind of perceptual content? This dissertation has defended an affirmative answer to this question. The difference between perception and thought lies in the different processes by which they generate states with representational content. Those differences result in states with different kinds of contents

My argument for this conclusion has proceeded through three main stages. Chapter two marked the first step, by defending a theoretically robust account of w*hat* it means for perception to represent the world 'nonconceptually' that was anchored in the structure of perceptual content. 'Conceptual contents' are structured propositions whose constituents are organized like a sentence: concepts, elements in such structures, refer to objects and properties in context-independent ways. We should think of the 'nonconceptual content of perception' as structured non-propositional contents whose constituents are organized like a graph or a property-space: the elements of perceptual contents refer to objects and properties in context-dependent ways; they achieve reference to mind-independent objects through their subjective effects on a subject's sense organs. Chapter three established a framework within which evidence could be brought to bear on this proposal, by showing how the Generality Constraint could be put to work as a negative criterion for nonconceptual content. As a positive criterion for conceptual content, the Generality

Constraint is justified by the role concepts play in explaining inferential capacities; capacities which use contents with freely recombinable constituents in the generation of new representational contents. As a negative criterion for nonconceptual content, the Generality Constraint invites us to look at restrictions on a subject's deployment of the constituents of their contents as possible evidence of non-inferential capacities to generate objective representations. Limited generality suggests a place for perception that elevates it above 'mindless' capacities for information-registration, without rising to the full objectivity of inferential thought. nonconceptual content is objective, but it is less objective than conceptual content. Chapter four gave substance to this outline sketch, by arguing for an account of perceptual capacities which substantiates the intuitive idea that perception is a 'receiving of the world.' Unlike the rational and conceptual schema governing thought, perceptual receptivity uses fixed and systematic organizing principles that proceed the registration of sensory information; the receiving of sensation is also, and already, a receiving of objective information. Whereas inferential capacities generate context-independent concepts, receptivity results in a context-dependent content, in which the real of sensation makes a semantic contribution to perceptual content.

The account I have defended engages the highly contentious 'nonconceptualism debate.' The view that perception has a 'nonconceptual content' has traditionally been opposed to a conceptualist view of experience. The most prominent representative of conceptualism in these pages has been John McDowell, although the view has also been defended by John Searle, Alex Byrne, Jerry Fodor and Zenon Pylyshyn. These philosophers acknowledge a theoretically deep difference between perception and thought, but deny

that this difference in kinds of mental states is accompanied by differences in mental contents. There is just one kind of content, organizing the world by means of *concepts* of objects and properties, to which subjects bear different relations. There are three main arguments that have been given for favoring a conceptual content over a nonconceptual content. First, some philosophers object to nonconceptual content on a priori grounds, since they regard all representation as essentially or constitutively concept-involving (Davidson, 1978; McDowell, 1994). Second, some philosophers object to nonconceptual content on the grounds that it cannot do justice to the epistemic status of perception. Perception justifies our knowledge of the empirical world. But if perceptual content is of a 'nonconceptual' kind, it is difficult to see how perception could stand in rational relations to beliefs, since rational relations hold amongst conceptually structured contents. Third, some philosophers argue that, whether or not it is coherent, the claim that perception has a nonconceptual content is simply unmotivated (Robbins, 2002; Byrne, 2004). Given Occam's razor, which says that the theory with the leaner ontology is to be preferred, conceptualism has the default status. The burden of proof lies with the nonconceptualist to identify some feature of perception which *could not* be captured save by appealing to a conceptual content—and they think that proponents of nonconceptual content have not met this burden of proof. Part of my argument for nonconceptual content has consisted in showing that these arguments are ungrounded. I think what drives both of the first two arguments is an assumption that to be 'nonconceptual,' perception would have to have a content that does not abstract or generalize or group particulars as belonging to types. This makes nonconceptual content seem incoherent because representational contents determine

conditions of correctness; and conditions of correctness fundamentally involve the idea of a way an object is presented as being—a way that the object can either be or fail to be (Brewer, 2006). And it makes nonconceptual content seem to lie outside of the sphere of justification because nonconceptual content lacks the kind of logical structure that could secure rational transitions between contents. Once we appreciate that the identification of classification with conceptualization is unwarranted, these objections collapse. Nonconceptual content has a 'spatial attributive' structure: perception classifies particulars as belonging to types; what makes it 'nonconceptual' is the fact that the attributive elements are not atomic. With respect to the argument from explanatory parsimony, meanwhile, the latter half of the dissertation challenges the conceptualist's claims to 'default status' on methodological grounds. Once we understand how the concept of content, and therefore of kinds of content is anchored in a scientific understanding of information-processing systems, this way of setting things up is revealed as ill-informed and unscientific. The best theory is the one that provides the most unified and elegant explanation for a given set of phenomena. We should look for a theory of perception, one which makes clear how all the parts—process, content, subject-matter—work together as a whole. Starting out with an assumption about what perceptual content must be like is inconsistent with a genuinely scientific approach to the mind.

The view that perception has a nonconceptual content view is also opposed to the views of naive realists, who share a commitment to the need for 'the nonconceptual' in perception, whilst holding a very different view about what this means. On the naive realist view, defended most recently by Charles Travis, Bill Brewer and John Campbell, seeing is

supposed to involve nothing more than a simple relation between the perceiver and the object seen-full stop. The relata of the relation are purely the perceiver and the objects or properties that are seen. Although ultimately I think we should reject this conception of the nonconceptual, one of my broader aims in this dissertation has been to show that the naive realist deserves much greater recognition for their contribution to the nonconceptualism debate. They are motivated by an idea of perception as fundamentally receptive, in contrast to the constructive and indirect nature of abstract, inferential thought—and they think that the structure of perceptual states must be concrete or immediate or particular to reflect this fact. I think they are right about both of these points. Where they go wrong is in supposing that 'concrete' and 'abstract' are mutually exclusive classifications—an assumption that is once again underwritten by the mistaken identification of classification with conceptualization. If we make this assumption, then the project is doomed—for the prospects of doing without representational content in explaining the operations of perceptual systems are slim. I suggest that nonconceptual content, understood in the way I propose, represents the only viable way to accommodate the naive realist's motivations whilst also doing justice to the facts uncovered by scientific psychology.

The arguments I have given moves this debate forward by uncovering and challenging the deeper assumptions motivating opposition to nonconceptual content that previous defenses of nonconceptual content have often overlooked. In this respect, the argument I have given for nonconceptual content is also critical of philosophers with whom I share an overarching commitment to nonconceptual content: Gareth Evans, Tim Crane, Fred Dretske, Richard Heck, Elizabeth Camp, and (to a lesser degree) Christopher Peacocke.

The lack of clarity that has surrounded conceptions of nonconceptual content since Evans is a failing for which Evans must bear a considerable amount of responsibility. It is unclear from Evans's discussion just what perceptual nonconceptual content is supposed to be. Is it a personal or subpersonal phenomenon? Is it concrete or partly abstract? Is it genuinely correctness-conditional or merely interpretable as such? Can it represent without supplement from conceptual capacities? Many philosophers who defend a 'nonconceptual content' have further contributed to these perplexities by not only failing to challenge but often actually *affirming* the banishment of nonconceptual content outside of the 'space of classification'—thus obscuring the rich possibilities inherent in nonconceptualist views of perception. In addition, existing defenses of nonconceptual content often damage the case for nonconceptual content through misguided methodological choices when arguing for its existence. Many nonconceptualists have focussed too much on the conceptualist, whilst failing to engage with and distinguish themselves from the naive realist who challenges nonconceptual content from the other side. They have often implicitly accepted the conceptualist's unwarranted demand for incontrovertible proof that perception could not be captured by a conceptual content; their failures to meet this impossible demand then lend weight to charges that nonconceptual content is unmotivated. I have also argued criticized nonconceptualists for the way in which they use the Generality Constraint as a negative criterion for nonconceptual content. These arguments rest on an empirically implausible view about what justifies our attribution of concepts to thinkers, fail to provide adequate grounds for distinguishing nonconceptual content from non-representational or mindless forms of engagement with the environment, and ignore the connections that exist

between a state's possession of content and the kind of contents that state has. In summary, the argument I offer in favor of nonconceptual content serves as much to motivate a different way of thinking about what nonconceptual content is and why we need it as it.

At the same time that my arguments do constitute genuine and substantive progress towards an understanding of the nature of perception and its relation to thought, they also leave open some important questions in need of further research. Since the present proposal remains essentially a framework sketch of perceptual capacities and contents, some of these take the form of demands for a filling-in of key details in this framework. Where do these perceptual spaces come from? Are they innate or acquired by perceivers in the early years of life? (And if they are innate, how did they come to be hardwired?) Can they vary amongst individuals? How does the general idea of a space play out within very different domains of magnitudes; for example, physical spaces as compared with color- or sound-spaces)? How do they operate within different perceptual modalities: what are the differences between visual spaces and those found in audition or touch? And how do these distinct property spaces—spaces for different kinds of properties or within different modalities—interact with each other in the generation of unified perceptual experiences? How can this framework accommodate attributives at higher levels—representations of facial expressions, representations of particulars as bodies—which scientific research strongly suggests are represented in perception.

Another set of questions concern how the present account bears on mental states beyond the case of perception. Are there mental states other than perceptions that represent the world nonconceptually—and, if so, to what extent do these states share

features in common with the nonconceptual content we find in perception? Do the two forms of classification I've argued for exhaust the forms which classification may take—or is classification a more robustly *pluralistic* (rather than dualistic) phenomenon? In perception, the nonconceptual character of content consists in the fact that perception represents properties in a way that is holistic, rather than atomic. This fact is in turn tied to the context-dependent or perspectival character of perceptual content, which is partly individuated by a subjective, sensory element. Certain cases of thought that use mental imagery seems to fit rather nicely with such an account. In *Thinking in Pictures*, Temple Grandin draws on her own experience as an autistic person, together with her extensive work with other autistic individuals, to argue that there are three different kinds of thinkers. Comparing her own 'picture-based' system of thought with the more widely-acknowledged linguistic kinds of abstraction, she writes,

> Growing up, I learned to convert abstract ideas into pictures as a way to understand them. … The Lord's Prayer was incomprehensible until I broke it down into specific visual images. The power and the glory were represented by a semicircular rainbow and an electrical tower. These childhood visual images are still triggered every time I hear the Lord's Prayer. The words "thy will be done" had no meaning when I was a child, and today the meaning is still vague. Will is a hard concept to visualize. When I think about it, I imagine God throwing a lightning bolt. Another adult with autism wrote that he visualized "Thou art in heaven" as God with an easel above the clouds. "Trespassing" was pictured as black and orange no trespassing signs. The word

"Amen" at the end of the prayer was a mystery: a man at the end made no sense.

(2006)

She goes on to describe the system by which she has learned to 'translate' between the two: when she hears somebody say the word "steeple," the first church that she sees in her imagination is almost always a childhood memory; but she has learned to form a generalized idea of the property of being a steeple by manipulating these specific images, imagining a church painted in different colors or put the steeple of one church onto the roof of another and not a church image that I have manipulated." This thinking in pictures constitutes a kind of abstraction closer to perception than to propositional thought. Recalling Kant's distinction between concepts and appearances—the one representing 'things in themselves,' the other representing objects 'insofar as they affect us'—we might say that what Grandin describes is a form of appearance-based thought. On the other hand, there are some kinds of primitive cognitive states, for example analog representations of number, which may be nonconceptual without having this imagistic or perspectival aspect. How are we to explain phenomena such as mental imagery and diagrammatic or pictoral inference? Especially given the emphasis I have placed on the variousness of mental representation, an acknowledgement of the vast array of psychological phenomena that do not fit neatly into these two domains is a crucial next frontier.

A number of years ago, I overheard a colleague (not a philosopher of mind) complaining that, whilst he had a feeling the nonconceptualism debate was a significant one, but none of the arguments he had heard make it clear to him why. If asked why we

should care about 'the nonconceptual content of perception,' one might respond by

pointing to the theoretical issues which turn on the existence of nonconceptual content:

"Issues about the individuation of conceptual content; about the nature of concept

possession; about the nature of rationality; about the relation between animal and human

perception; and even about our conception of objectivity all turn in part on the possibility

of nonconceptual content in perception" (Peacocke, 2001: 2). Yet my colleague's comment

was made after a talk I had given in which I had gone to great lengths to enumerate just the

sorts of theoretical issues Peacocke discusses. And if it troubled me greatly at the time, this

was because I felt I knew just what he meant. Here was an issue which I found sufficiently

intellectually exciting for it to have occupied my attentions for the better part of a decade.

Yet for all of their philosophical significance, the considerations adduced in this list are also

likely to bore the socks off anyone outside of the small circle of philosophers who work on

the theory of content.[44] In a similar vein, it had always frustrated me that I could not give

---

[44] 'Wandering Significance,' Mark Wilson's delightful book on concepts, begins with a (characteristically provocative) statement of the same sort of problem:

> to be honest, the central concerns of this book book—issues relating to the status of concepts, notions, properties, attributes, traits, characteristics and other notions of that ilk—have acquired a hard-won reputation for dullness, such that otherwise ardent students of philosophy frequently shun the subject as irrelevant to the normal run of human concerns. And the usual literature on the topic often confirms this somewhat leaden impression. (2006: 1)

The anecdote with which Wilson illustrates this point concerns a textbook which he received from a publisher, whose "fulsome blurb", praising its ability to make the subject of properties interesting to undergraduates, did not match its insides. A little box with the word ''the'' inscribed several times inside is accompanied by the question, ''How many 'thes' do you think are in the box?'' prompting the investigation of a lengthy sequence of theories of universals. Wilson's commentary on this (too enjoyable not to relate): "The enthusiast from the publicity department evidently believed that, in a classroom situation, some clever pupil will suggest the answer ''One'' and this startling proposal will ignite such heated debate that the entire class will sit in transfixed attention throughout an entire semester. For myself, I would not trust my pedagogy to such a slender motivational reed."

non-philosophically trained friends and family members even the vaguest intimations of what my work was about without recourse to convoluted summaries of theoretical issues about concepts and contents (of course, at the point when I started discussing the nature of Fregean propositions or Evans' Generality Constraint, the game was already lost). And it seems to me now, as it seemed to me back then, that we have failed to do justice to the topic of nonconceptual content if we have not communicated what Davidson once referred to as the 'heady exoticism' of that idea.

Borges's short story, 'Tlön, Uqbar, Orbis Tertius,' narrates the discoveries of a scholar concerning certain distant planets whose languages and intellectual cultures are unimaginably different from our own. In the language of Tlön's southern hemisphere, there are no nouns; only impersonal verbs, modified by monosyllabic suffixes (or prefixes) with an adverbial value: the analog for "moon" would be "to moonate"; the sentence, "The moon rose above the river" would be most accurately rendered as "upward behind the onstreaming it mooned." Their metaphysics is described as 'congenitally idealist': "the world for them is not a concourse of objects in space; it is a heterogeneous series of independent acts. It is successive and temporal, not spatial" (1962: 22). On Tlön's northern hemisphere, meanwhile, the prime unit is not the verb, but the monosyllabic adjective: "moon" would be translated as "round airy-light on dark" or "pale-orange-of-the-sky"; and "the metaphysics of this culture (like Meinong's subsistent world) abounds in ideal objects, which are convoked and dissolved in a moment, according to poetic needs" (ibid., 23).

I think the present distinction between conceptual and nonconceptual content—a difference in how the *world* is presented to subject, and not between one thing which is the

presentation of a world and something else that is sub- or pre-objective—yields a picture

of perception and thought as 'untranslatable languages' in something like Borges's sense.

Of course, the differences we find in perception do not go as far as those of Borges's

imagined worlds. Like conceptual thought, perception organizes the world into particulars

and properties ('nouns' and 'verbs'). Nevertheless, there are substantial differences within

these alternative systems concerning *what it is* to be a property. Percepts and concepts

both represent magnitudes: properties like being a given size or being a given color or

being oriented at a given angle. Concepts *reify* particulars and properties: a singular term is

(to use Frege's terminology) a 'saturated' kind of object; whilst a general term is an

'unsaturated kind of object,' standing in need of fulfilment by a singular term. But

perception does not reify properties: perceptual magnitudes are relationally individuated;

they are 'places in a system.' What it is to for something to possess such-and-such a length

(in perceptual terms) is for it to stand in certain relationship to the whole system of

lengths. One way in which these differences show up is in terms of our grasp of 'what there

is.' The perceptual world has a limited reach; it can represent only a limited subset of the

phenomena that we can think about. For example, perception can represent only those

domains of properties that form highly regular, reliable systems; and those particulars that

have a causal effect on the senses. But most of the properties we can think about are not

like that: you don't need to know all the species of dog to have the concept of a labrador. On

the other hand, there are also positive lessons to be learned from perception. As I

suggested in chapter two, the problem of the unity of the proposition plausibly arises from

the reifying turn we find in concepts; the problem looks very different when we approach it

from the point of view of spatial-attributive structure. We might think, too, about the expressive weaknesses of language; the struggle to achieve satisfactory expression of certain contents or phenomena which the poet T.S. Eliot memorably called, "the raid on the inarticulate." Although concepts open up certain possibilities to us that are not available to us in nonconceptual perception, the one is not 'better than' or more correct than the other. Taken together, they show the richness that is present in the way subjects engage with and think *about their worlds*.

REFERENCES

Allen, C., & Hauser, M. D. (1991). Concept attribution in nonhuman animals: Theoretical and methodological problems in ascribing complex mental processes. *Philosophy of Science*, 58(2), 221-240.

Anderson, J. (1978). Arguments Concerning Representations in Mental Imagery. *Psychological Review* 85(4): 249-277.

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

Beck, J. (2012). The Generality Constraint and the Structure of Thought. *Mind* 121(483): 563-600.

Beck, J. (2014). Analogue magnitude representations: A philosophical introduction. *The British Journal for the Philosophy of Science*, 66(4), 829-855.

Bell, D. (2001). The Inaugural Address: Some Kantian Thoughts on Propositional Unity. *In Aristotelian Society Supplementary Volume*. 75(1): 1-16. University College London: The Aristotelian Society.

Berkeley, G. (1709). An essay towards a new theory of vision. Dublin: Rhames & Papyat.

Berkeley, G. (1733). The theory of vision or visual language. London: J. Tonson.

Bermudez, J. L. (1995). Non-conceptual content: From perceptual experience to subpersonal computational states. *Mind and Language* 10(4): 333-369.

Bermudez, J. L. (2003). *Thinking Without Words*. New York, Oxford University Press.

Bigelow, J., Pargetter, R., & Armstrong, D. M. (1988). *Quantities. Philosophical Studies*, 54(3), 287-316.

Block, N. (2014). Seeing-as in the light of vision science. *Philosophy and Phenomenological Research*, 89(3), 560-572.

Borges, J. L. (1964). *Labyrinths: Selected stories & other writings.* New York: New Directions Publishing.

Brewer, B. (2006). Perception and content. *European Journal of Philosophy*, 14, 165–181

Brewer, Bill, (2000). *Perception and Reason*, Oxford: Oxford University Press.

Burge, T. (1995). Intentional Properties and Causation. In Burge, T. (2007). *Foundations of mind* (Vol. 2). New York: Oxford University Press.

Burge, T. (2005b). Disjunctivism and Perceptual Psychology. *Philosophical Topics*, 33, 1–78.

Burge, T. (2007). *Foundations of Mind. New York*: Oxford University Press.

Burge, T. (2009). Five theses on de re states and attitudes. In *The Philosophy of David Kaplan*. Almog, J. and Leonardi, P. (eds.). 246-324.

Burge, T. (2010). *Origins of Objectivity*. New York, Oxford University Press.

Burge, T. (2014). Reply to Rescorla and Peacocke: Perceptual content in Light of Perceptual Constancies and Biological Constraints. *Philosophy and Phenomenological Research* 88(2): 485-501.

Byrne, A. (2001). Intentionalism. *Philosophical Review 110* (April 2001): 199-240.

Byrne, A. (2003). Color and Similarity. *Philosophy and Phenomenological Research* 66(3): 641-665.

Byrne, A. (2005). Perception and Conceptual Content. In E. Sosa and M. Setup (eds), *Contemporary Debates in Epistemology.* Oxford: Blackwell, 2005, pp. 231–50.

Byrne, A. (2009). Experience and content. *The Philosophical Quarterly*, 59(236), 429-451.

Byrne, A. (2001). Intentionalism Defended. *Philosophical Review*, 110: 199–240.

Camp, E. (2004). The Generality Constraint and Categorical Restrictions. *The Philosophical Quarterly* 54(215): 209-231.

Camp, E. (2007). Thinking with Maps. *Philosophical Perspectives 21(Philosophy of Mind)*: 145-182.

Camp, E. (2009). Putting Thoughts to Work: Concepts, Systematicity, and Stimulus-Independence. *Philosophy and Phenomenological Research* 78(2): 275-311.

Campbell, J. (1994). *Past, Space, and Self.* Cambridge: MIT Press.

Carey, S. (2009). *The Origin of Concepts. New York*: Oxford University Press.

Carruthers, P. (2004): ''On being simple minded,'' American Philosophical Quarterly 41, 205–22; reprinted in Consciousness: Essays from a Higher-Order Perspective (Oxford: Oxford University Press, 2005).

Carruthers, P. (2006): The Architecture of the Mind: Massive Modularity and the Flexibility of Thought (Oxford: Oxford University Press).

Casati, R. V., Achille (1999). *Parts and Places: The structure of spatial representation*. Cambridge, MA, MIT Press.

Chalmers, D. (2006). Perception and the fall from Eden. In T. Gendler & J. Hawthorne (Eds.), *Perceptual experience* (pp. 49–125). Oxford: Clarendon Press.

Chuard, P. (2006). Demonstrative concepts without re-identification. *Philosophical Studies*, 130, 153–201.

Clapp, L. and Duhau. L. (2011). Varieties Of The Generality Constraint. *Manuscrito*, *34*(2), 397-434.

Clark, A. (2015). *Surfing uncertainty: Prediction, action, and the embodied mind*. New York: Oxford University Press.

Collins, J. (2011). *The unity of linguistic meaning*. Oxford: Oxford University Press.

Conan-Doyle, Arthur. (2001/1887). *A Study in Scarlet*. Hertfordshire: Wordsworth Editions.

Costall, A. P. (1984). Are theories of perception necessary? A review of Gibson's The ecological approach to visual perception. *Journal of the Experimental Analysis of behavior*, 41(1), 109-115.

Crane, T. (1988). The waterfall illusion. *Analysis*, 48, 142–147.

Crane, T. (1992). The nonconceptual content of experience. In T. Crane (Ed.), *The contents of experience: essays on perception* (pp. 136–157). Cambridge, MA: Cambridge University Press.

Crane, T. (1992a). Introduction. In T. Crane (Ed.), *The contents of experience: essays on perception* (pp. 1–17). Cambridge, MA: Cambridge University Press.

Crane, T. (2009). Is perception a propositional attitude? *The Philosophical Quarterly*, 59, 452–469

Crane, T. (2013). The Given. In Joseph Schear (ed.), (2013). Mind, Reason and Being-in-the-World: *The McDowell-Dreyfus Debate*. London: Routledge. pp. 229-249

Crane, T. and French, C. (2015). The problem of perception, *The Stanford Encyclopedia of Philosophy* (Spring 2016 Edition), Edward N. Zalta (ed.),

Crawford, D. (1982). Are There Mental Inferences in Direct Perceptions? *American Philosophical Quarterly*, 19(1), 83-91.

Crowther, T. (2006). Two conceptions of conceptualism and nonconceptualism. *Erkenntnis*, 65, 245–276.

Cummins, R. (1996). Systematicity. *The Journal of Philosophy*, 93(12), 591-614.

Cussins, A. (2003). Content, conceptual content, and nonconceptual content. In Y. Gunther (Ed.), *Essays on Nonconceptual Content* (pp. 133–163). Cambridge, MA: Bradford Books

Dasgupta, S. (2013). Absolutism vs Comparativism about Quantity. *Oxford Studies in Metaphysics*. K. B. a. D. Zimmerman, Oxford University Press. 8.

Dasgupta, S. (2014). On the Plurality of Grounds. *Philosophers Imprint* 14(20).

Davidson, D. (1973, January). On the Very Idea of a Conceptual Scheme. *In Proceedings and Addresses of the American Philosophical Association* (Vol. 47, pp. 5-20). American Philosophical Association.

Dennett, D. C. (1971). Intentional systems. Journal of Philosophy, 68, 87–106. Dennett, D. (1998). *Brainchildren: essays on designing minds*. London: Penguin Books

Dennett, D. C. (1991). *Consciousness Explained*. New York Back Bay Books / Little, Brown and Company.

Dickie, I. (2010). The generality of particular thought. *The Philosophical Quarterly*, 60(240), 508-531.

Dretske, F. (1969). Seeing and Knowing. Chicago: University of Chicago Press. Dretske, F. (1981). *Knowledge and the flow of information*. Cambridge, MA: MIT.

Dretske, F. (1981). Sensation and Perception. In *Knowledge and the Flow of Information*, Cambridge: The MIT Press, 135-153.

Dretske, F. (2006). Perception without Awareness. In T. Gendler & J. Hawthorne, (Eds.), *Perceptual experience*. Oxford University Press, 147-180.

Dreyfus, H. L. (2001). Introduction I: Todes. Todes, S. (2001). *Body and world*. MIT Press.

Duhau, L. (2009). Conceptuality and generality: a criticism of an argument for content dualism. Crtica, 41, 39–63.

Duhau, L. (2011). *Perceptual nonconceptualism: disentangling the debate between content and state nonconceptualism*. European Journal of Philosophy.

Eddon, M. (2013). Quantitative Properties. *Philosophy Compass* 8(7): 633-645.

Epstein, W. (1982). Percept-percept couplings. *Perception*. 11: 75-83

Evans, G. (1981). Semantic Theory and Tacit Knowledge. *In his* (1985), 322-42

Evans, G. (1982). *Varieties of Reference*. Oxford: Oxford University Press.

Evans, G. (1985). *Collected Papers*. Oxford: Clarendon Press.

Falkenstein, L. (1995). *Kant's Intuitionism*. Canada, University of Toronto Press Incorporated.

Field, H. (2001). Stalnacker on Intentionality. In his *Truth and the Absence of Fact*. Oxford: Oxford University Press.

Fine, K. (2007). *Semantic Relationism*. Oxford, UK, Blackwell Publishing.

Fine, K. (2010). *Towards A Theory of Part*. The Journal of Philosophy CVII(11): 559-589.

Fodor, F. A. (1983). *The Modularity of Mind*. Cambridge, Mass., MIT Press.

Fodor, J. (1975). *The Language of Thought*. New York, Thomas Y Crowell Company.

Fodor, J. (1987). *Psychosemantics*. Cambridge, MA: MIT.

Fodor, J. (2007). The Revenge of the Given. In , *Contemporary Debates in the Philosophy of Mind*. Ed. J. Cohen & B. McLaughlin.Oxford: Blackwell. pp.105–116.

Fodor, J. (2015). Burge on perception. In *The conceptual mind: New directions in the study of concepts*. Ed. Margolis, E. and Laurence, S. Cambridge: MIT Press. pp.203-222.

Fodor, J. and Pylyshyn, Z. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition* 28: 3-71.

Forge, J. (1995). Bigelow and Pargetter on quantities. *Australasian Journal of Philosophy* 73(4): 594-605.

Frege, G. (1893). *Grundgesetze der arithmetik* (Vol. 1). H. Pohle.

Frege, G. (1918/1956). The thought: A logical inquiry. *Mind*, 65(259), 289-311.

Frege, G. (1923/1963). Compound Thoughts. *Mind*, New Series, Vol. 72, No. 285, 1-17.

Gallistel, C. (1990). *The Organization of Learning* Cambridge, MIT.

Galistel, C. G., R (2000). Non-verbal numerical cognition from reals to integers. *Trends in Cognitive Sciences* 4(2): 59 - 65.

Gauker, C. (2011). Concepts are not icons. *Behavioral and Brain Sciences*, 34(3), 127-127.

Gauker, C. (2011). *Words and Images: An Essay on the Origin of Ideas*, Oxford University Press, 2011,

Gelman, R. G., CR (2004). Language and the origin of numerical concepts. *Science* 306: 441-443.

Gibson, J. J. (1960). The concept of the stimulus in psychology. *American psychologist*, 15(11), 694.

Gibson, J. J. (1963). The useful dimensions of sensitivity. *American psychologist*, 18(1), 1.

Gibson, J. J. (1966). The senses considered as perceptual systems. Boston: Houghton Mifflin.

Glüer, K. (2009). In Defence of a Doxastic Account of Experience. *Mind and Language*, 24: 297–327.

Gogel W. C. and Koslow M. (1972). The adjacency principle and induced movement. *Perception & Psychophysics* 11, 309-314

Gilchrist A. L. (1977). Perceived lightness depends on perceived spatial arrangement. *Science*. 195, 185-187

Grandin, T. (2006). *Thinking in pictures: And other reports from my life with autism*. New York: Vintage.

Gregory, R.L. (1974). Perceptions as hypotheses. In S.C. Brown (Ed.) Philosophy of psychology (pp. 195–210). London: Macmillan.

Hanna, R. (2000). Kantian non-conceptualism. *Philosophical Studies*, 137, 41–64.

Hanna, R. (2001). *Kant and the Foundations of Analytic Philosophy*, Oxford: Clarendon/Oxford University Press.

Hanna, R. (2005). Kant and Nonconceptual Content. *European Journal of Philosophy*, 13: 247–290.

Hanna, R. (2011). Kant's Non-Conceptualism, Rogue Objects, and the Gap in the B Deduction. *International Journal of Philosophical Studies*, 19: 399–415.

Hanna, R. (2016). Directions in Space, Nonconceptual Form, and the Foundations of Transcendental Idealism, in D. Schulting (ed.), *Kantian Nonconceptualism*, London: Palgrave Macmillan, pp. 99–115.

Heck, R. (2000). Non-conceptual Content and the Space of Reasons. *Philosophical Review*, (109): 483-523.

Heck, R. (2007). Are there different kinds of content?. In J. Cohen & B. McLaughlin (Eds.), *Contemporary debates in philosophy of mind* (pp. 117–138). Oxford: Blackwell

Hilbert, D. (2005). Color constancy and the complexity of color. *Philosophical Topics*, 33(1), 141-158.

Hilbert, D. (2012). Constancy, Content, and Inference. In Hatfield, G., & Allred, S. (Eds.). (2012). *Visual experience: Sensation, cognition, and constancy*. Oxford University Press. 199-211.

Hume, D. (1740/1975). *A Treatise of Human Nature*. Edited by L. A. Selby-Bigge, 2nd ed. revised by P. H. Nidditch, Oxford: Clarendon Press, 1975.

Hume, D. (1748/1975) An Enquiry concerning Human Understanding, contained in *Enquiries concerning Human Understanding and concerning the Principles of Morals*, edited by L. A. Selby-Bigge, 3rd ed. revised by P. H. Nidditch, Oxford: Clarendon Press, 1975.

Hurley, S. (1998). Vehicles, Contents, Conceptual Structure and Externalism. *Analysis* 58(1): 1-6.

Husserl, E. (1938/1973). *Experience and Judgment*. Trans. James S. Churchill and Karl Ameriks (Evanston, IL: Northwestern University Press).

Jacobs, Jane. (1961). *The Death and Life of the Great American Cities*. New York: Random House.

King, J. C. (2007). *The nature and structure of content*. Oxford: Oxford University Press.

King, J., and Speaks, J. (2014). *New Thinking about Propositions*. New York, Oxford University Press.

Johnson, K. (2004). On the systematicity of language and thought. *The Journal of Philosophy*, 101(3), 111-139.

Kanizsa, G. (1985). Seeing and thinking. *Acta psychologica*, 59(1), 23-33.

Kant, I. (1996/1787). *Critique of Pure Reason.* Translated by Werner S. Pluhar; introduced by Patricia Kitcher. Indiana: Hackett Publishing Company.

Kelly, S. (2001a). Demonstrative concepts and experience. *Philosophical Review*, 110, 397–420.

Kelly, S. (2001b). The non-conceptual content of perceptual experience: situation dependence and fineness of grain. *Philosophy and Phenomenological Research*, 62, 601–608.

King, J. (1996). Structured Propositions and Sentence Structure. *Journal of Philosophical Logic* 25(5): 495-521.

Locke, J. (1690/1975). An Essay Concerning Human Understanding: Edited with an Introduction, *Critical Apparatus and Glossary* by Peter H. Nidditch.

Marr, D. (1982). *Vision.* San Francisco: Freeman.

Martin, M. (1992). Perception, concepts, and memory. *The Philosophical Review*, 101, 745–763.

Martin, M. (1992). *Sight and Touch. In The Contents of Experience*. T. Crane. Cambridge, Cambridge University Press.

Martin, M.G.F. (2000). Beyond Dispute: Sense-Data, Intentionality and the Mind-Body Problem. *in Crane and Patterson* (eds.) 2000, 195–231.

Martin, M.G.F., (1992a). Perception, Concepts and Memory. *Philosophical Review*, 101:745–63.

McDowell, J. (1994). *Mind and World*. Cambridge, Mass: Harvard University Press.

McDowell, J. (1994a). The content of perceptual experience. *The Philosophical Quarterly*, 175, 190–205

McDowell, J. (2009d). The Logical Form of an Intuition. *In Having the world in view: essays on Kant, Hegel, and Sellars* (pp. 23–43). Cambridge, MA: Harvard University Press.

Medina, Jose. (2002). *Unity of Wittgenstein's Philosophy: Necessity, Intelligibility, and Normativity*, State University of New York Press, 2002.

Merleau-Ponty, Maurice, Phenomenology of Perception (1945/1966), trans. Colin Smith. London: Routledge & Kegan Paul.

Millikan, R. (1991). Perceptual Content and Fregean Myth. *Mind* 100(4): 439-459.

Millikan, R. G. (2012, November). Accidents. In Proceedings and Addresses of the American Philosophical Association(Vol. 86, No. 2, pp. 92-103). *American Philosophical Association.*

Millikan, Ruth Garrett. (1987) *Language, Thought and Other Biological Categories*, Cambridge, Mass.: MIT Press.

Mundy, B. (1988). Extensive Measuement and Ratio Functions. *Synthese* 75(1): 1-23.

Nagel, T. (1986). *The View From Nowhere*. New York: Oxford University Press.

Nietzsche, F. (1873/2006). On truth and lies in a nonmoral sense (1873). *The Nietzsche Reader*. Eds. Pearson, K. A., & Large, D. (2006). Oxford, Malden, Mass., and Carlton, Australia: Blackwell. pp. 114-123

Noë, Alva. (2004). *Action and Perception*, Cambridge, MA: MIT Press.

O'Keefe, J., and Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Oxford: Oxford

Palmer, S. (1997). Introduction to *Indirect perception*. Irvin Rock (author). Cambridge: MIT Press. Xi-xxvi.

Pargetter, J. B. a. R. (1988). Quantities. *Philosophical Studies* 54(3): 287-304.

Pautz, Adam, 2010, Why Explain Experience in Terms of Content? in Nanay (ed.) 2010. *Perceiving the world*. New York: Oxford University Press.

Peacocke, C. (1986). Analogue Content. Proceedings of the Aristotelian Society, *Supplementary Volumes* Vol 60: 1-17.

Peacocke, C. (1989). Perceptual Content. *Themes from Kaplan*. J. P. Joseph Almog, Howard Wettstein. New York, Oxford University Press: 297-330.

Peacocke, C. (1992). *A Study of Concepts*. Cambridge, Massachusetts, MIT Press.

Peacocke, C. (1994). Content, Computation and Externalism. *Mind and Language* 9(3): 303-335.

Peacocke, C. (1994). Non-conceptual content: kinds, rationales and relations. *Mind and Language,* 9, 419–429.

Peacocke, C. (1997). *Concepts without Words. Language, Thought and Logic*. R. Heck, Oxford University Press.

Peacocke, C. (1998). Nonconceptual content defended. *Philosophy and Phenomenological Research*, 58, 381–388.

Peacocke, C. (2001a). Does perception have a nonconceptual content?. *Journal of Philosophy*, 98, 239–264.

Peacocke, C. (2001b). Phenomenology and nonconceptual content. *Philosophy and Phenomenological Research*, 62, 609–615.

Peacocke, C. (2003). Postscript: the relations between conceptual and nonconceptual content. In Y. Gunther (Ed.), *Essays on nonconceptual content*. Cambridge, MA: Bradford Books. 318–322

Peacocke, C. (2014). Magnitudes: Metaphysics, Explanation and Perception. *Proceedings of the 2013,* Kirchberg Symposium.

Peacocke, C. (2015). Perception and the first person. In M. Matthen (Ed.), *The Oxford handbook of perception.* Oxford University Press.

Plotinus. (3rd century CE/2018). *The Enneads*, edited by Lloyd P. Gerson, and translated by George Boys-Stones, John M. Dillon, Lloyd P. Gerson, R.A. King, Andrew Smith and James Wilberding, Cambridge: Cambridge University Press.

Pylyshyn, Z. (1973). What the mind's eye tells the mind's brain. *Psychological Bulletin 80*: 1-24.

Pylyshyn, Z. (2007). *Things and Places: How the mind connects with the world*. Cambridge, MA, MIT Press.

Rescorla, M. (2009). Cognitive maps and the language of thought. The British Journal for the Philosophy of Science, 60(2): 377-407.

Rescorla, M. (2009). Predication and Cartographic Representation. *Synthese* 169(1): 175-200.

Rey, G. (1983). Concepts and stereotypes. *Cognition*, 15, 237-262

Rock, I. (1982). Inference in perception. In PSA: *Proceedings of the Biennial Meeting of the Philosophy of Science Association.* 2: 525-540.

Rock, I. (1983). *The logic of perception*. Cambridge MA: MIT Press.

Rock, I. (1985). Perception and knowledge. *Acta Psychologica*, 59(1), 3-22.

Rock, I. E. (1997). *Indirect perception*. Cambridge, MA.: MIT Press.

Rock, I. and Brosgole, L. (1964). Grouping based on phenomenal proximity. *Journal of Experimental Psychology*, 67(6), 531-538

Rosenthal, D. (2010). How to think about mental qualities. *Philosophical Issues* 20 (Philosophy of Mind).

Russell, B. (1919). On Propositions: What they are and How they Mean. *Proceedings of the Aristotelian Society*, Supplementary Volumes Vol. 2, *Problems of Science and Philosophy*, 1-43.

Russell, B. (1923). Vagueness. *The Australasian Journal of Psychology and Philosophy*, 1(2): 84-92.

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

Saint-Exupery, Antoine de. (2000/1943). *The Little Prince*. Translated by Richard Howard. New York: Harcourt.

Schellenberg, S. (2011). Perceptual content defended. *Noûs*, 45(4), 714-750.

Schulte, J. (2014). "We have a color system as we have a number system." In Wittgenstein on Colour. Ed. by Gierlinger, Frederik A. / Riegelnik, Štefan. Berlin: De Gruyter.

Shepard, R. N. (1984). Ecological constraints on internal representation: Resonant kinematics of perceiving, imagining, thinking, and dreaming. *Psychological review*, 91(4), 417.

Shephard., K. (2007). Spatial Processing in Navigation, *Imagery and Perception, Springer*.

Siegel, E. (1981). *Self and world: An explanation of aesthetic realism*. New York: Definition Press

Siegel, S. (2006). Which properties are represented in perception?. In T. Gendler & J. Hawthorne, (Eds.), *Perceptual experience*. Oxford: Oxford University Press. pp. 481–503

Siegel, S. (2012). *The contents of visual experience*. New York: Oxford University Press.

Siegel, S. (2013). *The contents of perception*. In E. N. Zalta (Ed.), The Stanford encyclopedia of philosophy (fall 2013 edition). Stanford University. http://plato.stanford.edu/archives/fall2013/ entries/perception-contents/

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

Soames, S. (2010). *Propositions.* In *The Routledge Companion to the Philosophy of Language*. Russell, G. and Graff-Fara, D. (eds.). New York: Routledge. pp. 209-220.

Speaks, J. (2005). Is there a problem about nonconceptual content?. *Philosophical Review*, 114, 359–398.

Stalnaker, R. (1998). What might nonconceptual content be? *Philosophical Issues*, 9: 339–352.

Stitch, S. (1983). *From folk psychology to cognitive science: the case against belief*. Cambridge, MA: MIT Press.

Strawson, P. (1959). *Individuals: an essay in descriptive metaphysics*. London: Methuen.

Swoyer, C. (1983). Realism and Explanation. *Philosophical Inquiry* 5: 14-28.

Swoyer, C. (1987). The metaphysics of measurement. *In Measurement, realism and objectivity* . Springer, Dordrecht. 235-290.

Todes, S. (2001). *Body and world*. MIT Press.

Toribio, J. (2008). State Versus Content: The Unfair Trial of Perceptual Nonconceptualism. *Erkenntnis*, 69: 351–361.

Toribio, J. (2011). Compositionality, iconicity, and perceptual nonconceptualism. *Philosophical Psychology*, 24: 177–193.

Travis, C. (1994, January). On constraints of generality. *In Proceedings of the Aristotelian Society* (Vol. 94, pp. 165-188). Aristotelian Society, Wiley.

Travis, C. (2004). The silence of the senses. *Mind*, 113(449): 57-94.

Travis, C. (2007). Reason's reach. *European Journal of Philosophy*, 15(2), 225-248.

Tye, M. (2005). On the nonconceptual content of experience. In M. Reicher & J. Marek (Eds.), Experience and analysis: *proceedings of the 27th international Wittgenstein symposium* (pp. 221–242). Vienna: öbv&hpt.

Tye, M. (2006). Nonconceptual content, richness, and fineness of grain. In T. Gendler & J. Hawthorne (Eds.), *Perceptual experience* (pp. 504–530). Oxford: Oxford University Press.

Walsh, V. (2003). A Theory of Magnitude: Common cortical metrices of Time, Space and Quantity. *Trends in Cognitive Sciences*, 7, 483-8.

Wilson, M. (2008). *Wandering significance: An essay on conceptual behavior*. New York: Oxford University Press.

Wittgenstein, L. (1922). *Tractatus Logico-Philosophicus*. London: Routledge.

Wittgenstein, L. (1929). 'Some Remarks on Logical Form'. *Proceedings of the Aristotelian Society, Supplementary Volumes*, *9*, 162-171.

Wittgenstein, L. (1953 (2001)). *Philosophical Investigations*. Oxford, Blackwell Publishing.

Wittgenstein, L. (1977). *Remarks on Colour*. Oxford, Blackwell Publishing.