

Talking and Knowing about How Things Look

Wylie Breckenridge

B.Phil. Thesis

February 2003

Abstract

In the first two chapters, I develop a theory of the truth conditions of sentences involving the word 'looks', such as 'Her shirt looks red' and 'That watch looks to be expensive'. In Chapter 1 I argue that to make sense of our 'looks' talk we must take it (i) that there are such things as ways of looking, (ii) that when objects look a certain way it is in virtue of having a way of looking, and (iii) that when we talk about the way an object looks we are talking about the way of looking that it has. I make comments about the identity conditions of ways of looking, about what kinds of thing they are, and about how we experience them.

In Chapter 2, I argue that we use 'looks' sentences to express propositions of three distinct forms, in some ways agreeing and in some ways disagreeing with Frank Jackson's well known account. I point out that many of our 'looks' sentences are ambiguous, stress the need to be clearer about which propositions we intend or understand them to express, and propose methods for disambiguation. In particular, I argue that 'looks the same as' can be used to express several relations, some of which are transitive and some of which are not, and for some of which it is true that if x looks the same as y and x looks F it follows that y looks F , and for some of which it is not. Finally, I argue that it follows from the theory developed in these first two chapters that a visual experience of an object does not represent the object to be a certain way, and that it cannot have the scenario content that Christopher Peacocke suggests it does. I suggest why it might be plausible to think so, offer an alternative account of what it is for the experience to be veridical, and propose that ways of looking might be a more appropriate basis for nonconceptual content theories of visual experience.

In the third and final chapter, I defend the theory against an apparently devastating charge: that it implies (a) that 'looks the same as' does not mean 'is perceptually indiscriminable from', and (b) that we have no conceptual guarantee that we can tell how things look to us just by looking at them. I give independent reasons for thinking that (a) and (b) are both true, offer an explanation for why we might nevertheless find them intuitively unattractive, and show that by accepting (b) we can offer an epistemic solution to phenomenal sorites paradoxes of the kind that have been offered for non-phenomenal ones, suggesting that each is a distinct instance of a more general epistemic phenomenon.

Acknowledgments

For many helpful discussions, comments and suggestions, thanks to Timothy Williamson, Richard Price, Bill Brewer, Sarah Moss, Nico Silins, and to audiences of The Ockham Society and the New College Graduate Colloquium to whom I presented sections of this thesis.

Contents

1	Ways of Looking	5
1.1	Ways	5
1.2	Having Ways	7
1.3	The Way Something Has	8
1.4	Respects	10
1.5	Talking about Ways	11
1.6	Identity of Ways	13
1.7	The Way Something Looks	16
2	Talking About How Things Look	22
2.1	The Theory	23
2.2	One Way of Talking	23
2.3	A Second Way of Talking	28
2.4	A Third Way of Talking	33
2.5	The Three Ways	36
2.6	Two Results	38
2.6.1	The Transitivity of ‘Looks the Same As’	38
2.6.2	A Truism about ‘Looks the Same As’ and ‘Looks F’?	40
2.7	Visual Experience and its Content	41
3	Knowing About How Things Look	45
3.1	Some Plausible Principles	45
3.2	Problems for Principle1	50
3.2.1	Mismatch with Other Ways	50
3.2.2	Exactly?	51
3.2.3	The Muller-Lyer Illusion	52
3.3	The Plausibility of Principle1 and Observationality	55
3.4	The Problem of the Non-Transitivity of <i>Matching</i>	58
3.5	Phenomenal and Non-Phenomenal Sorites Series	63
	Bibliography	68

Chapter 1

Ways of Looking

In the first two chapters I will develop a theory of the truth conditions of sentences involving the word ‘looks’, not when it is used in the sense in which a man looks at his wife when she is talking to him, but in the sense in which she looks pleased when he does. The sentences that I want to fall within the scope of the theory include ‘Her shirt looks red’, ‘That watch looks to be expensive’, ‘The players look like ants’, ‘The top line looks longer the bottom line’, and ‘The first patch looks the same colour as the second one’. Intuitively, they are the sentences that we use to talk about the way something looks, when ‘looks’ is being used in the intended sense. Syntactically, they are sentences of the form ‘x looks ...’ where ‘x’ is a singular referring expression.

I will do so by considering the way that we talk about ways of looking, and assimilating such talk to the way that we talk about ways more generally: ways of walking, ways of thinking, ways of being, ways of making the most of one’s opportunities, and so on. I will argue that to make sense of such talk we need to take much of it to involve quantification over ways, and to be a matter of describing and comparing them. In particular, to make sense of the way we talk about the way things look we need to take much of it to involve quantification over, description of, and comparison of ways of looking.

1.1 Ways

We talk as if there are ways over which we can quantify. A chef says to his apprentice that there is a way of chopping onions that produces fewer tears; a geography teacher tells his students that there are many ways of saving water; a bank robber says that of all the ways of making money his is the most fun; and someone once said that where there’s a will there’s a way.

In particular, we talk as if there are ways of being and ways of looking over which we can quantify. I say that there is a way that I want to be when I grow up, and you ask me which way that is; I ask you how (in which way) you would like to be, and you say that any way that makes you happy will do; a photographer tells her model that there is a way of looking that she thinks he should try, and the model asks the photographer which way that is; a printer asks the customer how (in which way) he would like the invitation to look, and the customer replies that any of those ways will do (pointing to the catalogue).

This talk commits us to there being ways. When I say, for example, that there is a way I want to be when I grow up, it follows from what I say that there are ways of being: it would be inconsistent of me to say that there is a way I want to be but that there are no ways of being. But it does not commit us to ways being things of any particular kind. It follows from what I say in this example that there are ways of being. It does *not* follow that these ways of being are concrete objects; they may be abstract objects: bundles of properties, classes of individuals, or something else. All that follows is that *there are* ways of being, whatever they might be. Nor does it commit us to this being the only way that we can talk. We do talk as if there are ways, but perhaps we could get by without doing so – perhaps by paraphrasing to remove all such talk, just as we might paraphrase what we say about average families to remove all talk that commits us to such things. But that would be to talk a different way. We are committed to there being ways in the sense that we do talk as if there are ways, even though we might not be committed to being so committed (i.e. even though we might not have to talk that way).

I am interested here in the way that we *do* talk. We *do* talk as if there are ways that we can quantify over, and we take ourselves to be doing just that in each of the examples above. So I will take it that when the chef says to his apprentice, for example, that there is a way of chopping onions that produces fewer tears, he is expressing a proposition of the form $(\exists w)\Phi(w)$, where the variable ‘w’ ranges over ways of chopping onions; and I will take it that when the bank robber says, for example, that of all the ways of making money his is the most fun, he is expressing a proposition of the form $(\forall w)\Phi(w)$, where ‘w’ ranges over ways of making money.

Whether or not we accept that there are possible worlds, and, if we do, whatever we think they are, taking there to be such things over which we can quantify allows us a powerful account of the semantics of our modal talk. In the same way, whether or not we accept that there are ways of looking, and, if we do, whatever we think they are, taking there to be such things over which we can quantify allows us, I hope to show, a powerful account of the semantics of our ‘looks’ talk.¹

1.2 Having Ways

Whatever ways are, we talk as if things can *have* them. An ecologist remarks that nature has an effective way of reducing excess population; I note that you have a way of getting things done which makes me look lazy; a friend remarks that my father and I have ways of walking that are very similar; but I add that we have ways of talking that are quite different. In particular, we talk as if things can have ways of being and ways of looking. I say to my off-beat friend that she has a way of being that makes me feel conservative; a woman is accused of having a way of being that makes men fall for her; I remark to my friend that she has a way of looking that drives me crazy, and she asks how I feel about the other ways of looking that she has.

This talk commits us to things being able to have ways, that is, to there being a relation of *having* that can obtain between things and ways. When I say that you have a way of getting things done that makes me look lazy, it follows from what I say that people can have ways of getting things done: it would be inconsistent of me to say that you have a certain way of getting things done but that people cannot have ways of getting things done. But it does not commit us to this relation of *having* being a relation of any particular kind. People can have ways of getting things done, but perhaps not in the same way that they can have names, hands, or jealous inclinations. And, again, it does not commit us to this being the only way that we can talk. We *do* talk as if things can have ways, but perhaps we could get by without doing so. But that would be to talk a different way. We are committed to things being able to have ways in the sense that

¹ John Knox Jr. (1967) argues that statements about the ways in which objects appear entail the existence of ways of appearing. Jason Stanley and Timothy Williamson (2001) quantify over ways, claiming that ‘... any successful account of natural language must postulate [such] entities’ (p. 427). Donald Davidson

we do talk as if things can have ways, even though we might not be committed to being so committed.

So I will take it that when I say to my friend, for example, that she has a certain way of looking (or that she looks a certain way) I am expressing a proposition of the form $(\exists w)\text{Has}(a,w)$, where ‘Has’ stands for this relation of *having* that can obtain between objects and ways of looking.

1.3 The Way Something Has

If we talk as if there are ways then we shouldn’t be surprised to find that we talk as if referring to ways. And it seems that we do. You ask me how I’d like you to look, and I say ‘*that way*’, pointing to a photo in a magazine, thereby using demonstration to refer to a way of looking. You ask me which is the best way to chop onions, and I say ‘the way that Peter does’, thereby using a definite description to refer to a way of chopping onions. I will take it that when we talk in this way we are referring to, or at least purporting to refer to, ways.

I might refer to the way that my friend walks, or the way that she used to walk, or the way that she would have walked had she been taller, or the way that she walks on Mondays, or the way that she walks in respect of stride-length. Am I referring, or purporting to refer, to a way of walking in each case? In those that I am, are those ways of walking distinct, or are some of them identical (even if different in some way)? In those that I am not, to what kind of thing am I referring, or purporting to refer?

Consider how we might answer similar questions about names. I might refer to the name of my friend, or the name that she used to have (suppose that she is married and has changed her name), or the name that she would have had had she been male, or the name that she has on weekdays (suppose that she uses her maiden name at work), or her first name, or her last name. I think that the following are reasonable things to say about this. The name that she used to have, ‘Kerrie Hawkins’, is distinct from the name that she now has, ‘Kerrie Cook’. It is not that there is a name, n , such that (i) n was and is my friend’s name, (ii) n used to be ‘Kerrie Hawkins’, but (iii) n is now ‘Kerrie Cook’.

(1967) argues that an analysis of action sentences ought to involve quantification over actions, in a way that I suggest can be adapted to suit the case of ‘looks’ sentences and ways of looking.

Rather, there are distinct names n_1 and n_2 such that (i) n_1 was and is ‘Kerrie Hawkins’, (ii) n_2 was and is ‘Kerrie Cook’, and (iii) n_1 was my friend’s name but now n_2 is my friend’s name. Similarly, the name that she would have had now had she been male, ‘Kevin Cook’, is distinct from the name that she does have now, ‘Kerry Cook’; the name that she has on weekdays, ‘Kerrie Hawkins’ is distinct from the name that she has on weekends, ‘Kerrie Cook’; and her first name, ‘Kerrie’, is distinct from her last name, ‘Cook’, and both are distinct from her full name ‘Kerrie Cook’ (all three are names, but the latter is some kind of structured complex that has the first two as constituents).

We can put these comments in terms of functions. There is, according to this account, no function f^1 such that if x is a person and t is a time then $f^1(x)$ is the name of x at t , because the one person may have different names at different times. Nor is there a function f^2 such that if x is a person, t is a time and Ω is a possible world then $f^2(x,t)$ is the name of x at t in Ω , because the one person at the one time may have different names in different possible worlds. But there is a function f^3 such that $f^3(x,t,\Omega)$ is the name of x at t in Ω , because the one person at the one time in the one possible world has only one name. But this is so only if we take the term ‘name’ to refer to a *full* name (like ‘Kerrie Cook’) and not parts of full names (like ‘Kerrie’ and ‘Cook’). Then we can say that the one person at the one time in the one world *does* have just one name, which may itself be constituted by a first name and a last name.² We can, if we like, take there to be a function f^3_1 such that $f^3_1(x,t,\Omega)$ is the first name of x at t in Ω , and a function f^3_2 such that $f^3_2(x,t,\Omega)$ is the last name of x at t in Ω . There may be something arbitrary about doing so, but it may be expedient for reasons of definiteness and clarity.

I think that similar remarks can be made about ways of walking. When the term ‘way of walking’ is taken to refer to the appropriate kind of thing, there is a function *Walks* such that if x is a person, t is a time, and Ω is a possible world then *Walks*(x,t,Ω) is the way that x walks at t in Ω , because the one person at the one time in the one possible world has only one way of walking. It might be thought that this is not so, because a person can have more than one way of walking at the one time in the one world: she might have a way of walking in respect of stride-length, a way of walking in respect of arm-swing, a way of walking in respect of head movement, and so on. But

² Strictly speaking this is false. Elton John (Reginald Kenneth Dwight) has more than one full name now in this world. Let me make the appropriate idealising assumption for the sake of this example.

that would be to misunderstand the intended sense of ‘way of walking’. In the intended sense, the term ‘way of walking’ is reserved for things like the way my friend walks, simpliciter, unqualified by any respect, and terms of the form ‘way of walking in respect of R’ for ways of a different kind. Then it is true to say that the one person at the one time in the one world *does* have just one way of walking, which may itself be constituted by a way of walking in respect of stride-length and a way of walking in respect of arm-swing, and so on. And we can, if we like, take there to be a function $Walks_1$ such that $Walks_1(x,t,\Omega)$ is the way that x walks in respect of stride-length at t in Ω , and a function $Walks_2$ such that $Walks_2(x,t,\Omega)$ is the way that x walks in respect of arm-swing at t in Ω , and so on. There may be something arbitrary about doing so, but it may be expedient for reasons of definiteness and clarity.

And I think that similar remarks can be made about all other ways, and in particular ways of being and ways of looking. So I will take it, for definiteness and clarity, that ways of being are such that there is a function Is such that if x is a person, t is a time, and Ω is a possible world then $Is(x,t,\Omega)$ is the way that x is at t in Ω , because the one person at the one time in the one possible world has only one way of being. These ways of being are structured complexes whose constituents include ways of being in certain respects (in respect of age, in respect of height, in respect of colour, and so on). And I will take it, for definiteness and clarity, and for the time being,³ that ways of looking are such that there is a function $Looks$ such that if x is a person, t is a time, and Ω is a possible world then $Looks(x,t,\Omega)$ is the way that x looks at t in Ω , because the one person at the one time in the one possible world has only one way of looking. These ways of looking are structured complexes whose constituents include ways of looking in certain respects (in respect of colour, in respect of shape, in respect of length, and so on).

1.4 Respects

What are these ‘respects’ that we seem to be referring to in expressions like ‘the way she walks in respect of stride-length’, ‘the way it is in respect of age’, and ‘the way it looks in respect of colour’?

We talk as if there are such things, and that we can quantify over them. A teacher tells her students that there is a respect in which a red square and a blue square are the same, and a respect in which they are different. You ask me to compare the way my favourite cup looks with the way that yours looks, and I say that in some respects they are quite similar but in most respects they are very different. A customer asks the salesman in which respect this model is better than that one, and the salesman says that it's better in every respect.

This talk commits us to there being respects. If there is a respect in which a red square and a blue square are the same then it follows that there are respects: it would be inconsistent to claim that there is such a respect but that there are no respects. It does not commit us to there being respects which are things in the way that physical objects are things, and it may not commit us to this being the only way that we can talk. But we do talk as if there are respects, and such talk commits us to them. So I shall take it that to specify what we mean by some of what we say requires quantification over respects.

We take it that each of the following is a respect: colour, shape, size, temperature, texture, moistness, pitch, clarity, stride-length, and arm-swing. What is a respect? Perhaps something like a degree of freedom in which ways can differ. The way that two things look can differ in respect of colour, shape, size, and so on. The way that two things feel can differ in respect of temperature, texture, moistness, and so on. The way that two people talk can differ in respect of pitch, clarity, hesitation, and so on. I will take it that respects are classes. The elements of the respect *colour* are colours: red, green, blue, and so on. The elements of the respect *shape* are shapes: square, circular, rectangular, and so on. And I will take it that some respects have structure. There is, for example, a containment relation on colours in virtue of which it makes sense to say that a certain two colours are both shades of the colour red, and a partial ordering on shapes in virtue of which it makes sense to say that one rectangle is closer to square than another. The significance of such structures will become evident in Section 2.4.

1.5 Talking about Ways

³ I will take this back somewhat in Section 1.7.

If we talk as if there are ways, then we shouldn't be surprised to find that we talk as if *about* ways. And it seems that we do. You ask me to tell you about the way my grandmother is and I tell you that she is old but smart and determined. You ask me to tell you about the way that she walks, and I tell you that she walks slowly, a bit unsteadily, but more confidently than most people her age. You ask me to tell you about the way that she makes scones, and I tell you that she simply mixes three parts flour with one part cream and one part lemonade and bakes them in a hot oven for fifteen minutes. In each case not only do I take myself to be talking about my grandmother, but I take myself to be talking about the way she is, the way she walks and the way she makes scones.

Sometimes we talk about ways by ascribing them properties, both non-relational and relational: the way that my grandmother makes scones *is* very easy, *is* easier than the way that my aunty makes scones; the way that some people look *is* very appealing, *is* more appealing than the way that I look.

But some of what we say about ways we only say by ascribing properties, non-relational and relational, to things that have those ways. When talking about the way my grandmother walks I don't say that it *is* slow, but that she walks slowly; and I don't say that the way she walks *is* slower than the way that I walk, but that she walks more slowly than I do. Ways of walking, we take it, cannot be slow, nor slower than other ways of walking; rather, people can walk slowly, more slowly than others.⁴ To describe the way that my favourite cup looks I don't say that the way it looks *is* green or that the way it looks *is* shiny but that the cup looks green and that the cup looks shiny. Ways of looking, we take it, cannot be green and ways of looking cannot be shiny; rather, things can look green and things can look shiny.⁵ It seems natural to say that I am talking about my grandmother by ascribing properties to her (the properties of walking slowly, and of walking more slowly than me) and about my cup by ascribing properties to it (the properties of looking green, and of looking shiny). And it seems natural to say that I am thereby talking about the way my grandmother walks and about the way my cup looks. But it does not seem natural to say that I am talking about the way my grandmother

⁴ I might, in apparent contradiction to this, say that she has *a slow way of walking*. But because this can naturally be taken as saying that she walks slowly, but not naturally taken as saying that the way she walks is slow, I suggest that it is just a way of saying the former rather than the latter.

⁵ Again, I might, in apparent contradiction to this, say that the cup has *a green way of looking*. I would make similar remarks about this case too.

walks or about the way my cup looks by ascribing properties to those ways. It seems that we are talking about these ways, but not by ascribing them properties.

There is a natural explanation for this. We ascribe a property to something by saying something about how it is, or about how it looks, or about how it feels, or about how it walks, or about how it lives, and so on. If there is no way that it is, or at least no way that it is in a certain respect, then we can't talk about the way that it is, or at least about the way that it is in that respect. If there is no way that it looks, or at least no way that it looks in a certain respect, then we can't talk about the way that it looks, or at least about the way that it looks in that respect. If there is no way that it has at all, then we can't talk about any way that it has, and so we cannot ascribe it any properties at all. Or, less drastically, if there is some way that it does not have, or at least some respect of some way that it does not have, then we can't talk about that way that it has, or at least about that way that it has in that respect, and so we cannot ascribe to it any properties that require us to talk about such things.

The natural explanation, then, for why for a given way, *w*, there are some properties that we can ascribe to *w* but some that we cannot, is that there are some ways which *w* does not have, or at least some respects of some ways that *w* does not have. The way that my grandmother makes scones has, itself, a way of being, and in particular a way of being in respect of easiness, and it is in virtue of it having this way that we can say that the way that my grandmother makes scones is very easy. The ways that people look have, themselves, ways of being, and in particular ways of being in respect of appeal, and it is in virtue of them having these ways that we can say that the way that some people look is more appealing than the way that I look. But the way that my cup looks has, itself, no way of being in respect of colour or in respect of shine, so we cannot say that the way my cup looks is green (or any other colour), or that the way my cup looks is shiny (or any other shininess). It is not that anything *is* green (in particular, it is not that any way of looking is green), just that something *looks* green.

1.6 Identity of Ways

If we talk as if there are ways and as if about ways, then we shouldn't be surprised to find that we talk as if comparing ways. And it seems that we do. I say that Bill runs

more smoothly than Ben, taking myself to have compared the way that Bill runs with the way that Ben runs. I say that this box is bigger than that box, taking myself to have compared the way that this box is with the way that that box is. I say that patch A looks more round than patch B, taking myself to have compared the way that patch A looks with the way that patch B looks.

In particular, we shouldn't be surprised to find that we talk as if *identifying* ways. I am watching Bill and Ben walk, and notice that in many respects the way that Bill walks is the same as the way that Ben walks. I comment to my friend that Bill walks the same as Ben. She too is watching them walk, and disagrees with me, pointing out that Bill swings his arm higher than Ben does. I agree, conceding that in *that* respect the way Bill walks is different from the way Ben walks, commenting more accurately that in respect of stride-length and posture and smoothness it is the same. And she agrees: Bill does not walk the same as Ben simpliciter, but he does walk the same in many respects.

Bill now starts to mimic the way that Ben walks, swinging his arms not quite as high and making every other necessary adjustment to the way that he walks. Now my friend and I agree that Bill walks the same as Ben in every respect, and that we can truly say that Bill walks the same as Ben, or that the way Bill walks is the same as the way that Ben walks, or that Bill and Ben walk the same way.

But a third friend joins us in watching Bill and Ben walk, and disagrees when we say that Bill and Ben walk the same way. He points out that they are walking differently in respect of which side of the footpath they walk on: Bill is walking on the left side, but Ben is walking on the right. We agree, but maintain that the way Bill walks is nevertheless the same as the way Ben walks, claiming that a difference in respect of which side of the footpath they walk on is not a relevant difference. In fact, we claim that the way that Bill walks is *identical to* the way that Ben walks, despite this difference.

What is going on here? Our mutual friend is not disagreeing with us about which are the respects in which Bill walks the same as Ben and which are those in which he walks differently: For every respect R, she agrees with us about whether or not Bill walks the same as Ben in respect of R. She is disagreeing with us about whether or not the respects in which Bill walks the same as Ben are sufficient to declare truly that the way Bill walks is identical to the way Ben walks. That is, she is disagreeing with us

over the *identity conditions* of ways of walking. She thinks it is a necessary condition for the way that Bill walks to be identical to the way that Ben walks that Bill walks the same as Ben in respect of which side of the footpath they walk on. We don't.

The important point is this: even though we disagree which class it is, we agree that there is *some* class of respects, call it $\mathfrak{R}_{\text{walking}}$, such that x walks the same as y if and only if for all respects R in $\mathfrak{R}_{\text{walking}}$, x walks the same as y in respect of R. I will call this the class of *respects of ways of walking*.

I could make similar remarks about other ways, and in particular about ways of being and ways of looking. We take it that there is some class of respects, call it $\mathfrak{R}_{\text{being}}$, such that x *is* the same as y if and only if for all respects R in $\mathfrak{R}_{\text{being}}$, x is the same as y in respect of R. The class $\mathfrak{R}_{\text{being}}$ probably includes such respects as size, shape, and colour, but probably does not include location. And we take it that there is some class of respects, call it $\mathfrak{R}_{\text{looking}}$, such that x *looks* the same as y if and only if for all respects R in $\mathfrak{R}_{\text{looking}}$, x looks the same as y in respect of R. The class $\mathfrak{R}_{\text{looking}}$ probably also includes such respects as size, shape, and colour, but probably not location. Following the above, I will call the members of $\mathfrak{R}_{\text{being}}$ the *respects of ways of being* and the members of $\mathfrak{R}_{\text{looking}}$ the *respects of ways of looking*.

There may be a correlation between the respects of ways and what we normally call *intrinsic properties*. If there is, I suggest it is this: The property of being r in respect of R is an intrinsic property of things if and only if R is a respect of ways of being; the property of looking r in respect of R is an intrinsic property of things if and only if R is a respect of ways of looking; the property of walking r in respect of R is an intrinsic property of things if and only if R is a respect of ways of walking. In general, the property of W-ing r in respect of R is an intrinsic property of things if and only if R is a respect of ways of W-ing. I suggest, also, that there is the following correlation between ways of looking and what are sometimes called *visual observational predicates* ('VOP' for 'visual observational predicate'):

- (VOPa) The predicate 'G' is a visual observational predicate if and only if, for all objects x and y, if the way that x looks is identical to the way that y looks then if 'G' applies to x then 'G' applies to y;

from which it follows:

- (VOPb) The predicate ‘looks F’ is a visual observational predicate if and only if ‘F’ names an element of some respect of ways of looking.⁶⁷

1.7 The Way Something Looks

I said in section 1.3 that I will take ways of looking to be such that the one object at the one time in the one possible world has only one of them, so that there is a well-defined function *Looks* such that if *x* is an object, *t* is a time, and Ω is a possible world then *Looks*(*x*,*t*, Ω) is the way that *x* looks at *t* in Ω . I now want to take some of that back.

We seem to talk about ways of looking as though the one object at the one time in the one possible world can have more than one of them. Bill and Ben are standing at different locations in a room, looking at a cylindrical object in the centre. They are asked to describe the way the cylinder looks. Bill can only see the cylinder end-on, and reports that from where he is it looks circular. Ben can only see the cylinder side-on, and reports that from where he is it looks rectangular. If Bill has said that the cylinder looks circular and Ben has said that the cylinder looks rectangular, then they ought to take themselves to have made inconsistent claims, and therefore to be in disagreement. For suppose that Bill and Ben now come together and, from the same location (near enough), look at a colour patch on the wall. Bill says that it looks circular, Ben says that it looks rectangular. They enter into an argument, on the basis that they can’t both be right. The patch, they take it, can’t both look circular and look rectangular – to say so would be to say inconsistent things about the way the patch looks. Or to see this another way, Bill and Ben agree that from where Bill is the cylinder cannot look both circular and rectangular: If Bill were to say so, then they both would take Bill to be making an inconsistent claim. Similarly, they agree that from where Ben is the cylinder cannot look both rectangular and circular: If Ben were to say so, then they both would take Ben to be making an inconsistent claim.

⁶ I don’t want to say that ‘looks F’ is a visual observational predicate just in case ‘F’ is *any* adjectival expression, because I think that’s false. See Sections 2.3 and 2.6.2.

⁷ Similar principles could be formulated for the correlation between ways of sounding and aural observational predicates, between ways of tasting and gustatory observational predicates, and so on.

But Bill and Ben don't take themselves to be in disagreement. Bill agrees with Ben that from where Ben is the cylinder looks rectangular, and Ben agrees with Bill that from where Bill is it looks circular, and they agree that the way the cylinder looks from one location is different from the way it looks from the other. So they must take it that they are talking about *distinct* things, things that they refer to as 'the way the cylinder looks from where Bill is' and 'the way the cylinder looks from where Ben is'. If we take these things to be ways of looking, then it seems that the one object at the one time in the one possible world can have more than one way of looking.

What should we say about this? Consider the way we talk about the ways things *are*, rather than the way they look. We normally take it that the properties of being circular in shape and being rectangular in shape are incompatible - the one object cannot be both circular and rectangular at the one time. If I said to you that I was holding a piece of cardboard that was both circular and rectangular, then you would no doubt take what I said to be false. But sometimes we seem to accept that the one object can be, and in fact is, both circular and rectangular. What shape is the cylinder that Bill and Ben are looking at? In one sense it is circular - the intersection of the cylinder with a plane parallel to its ends and midway between them is a surface circular in shape. But in another sense it is rectangular - the intersection of the cylinder with a plane perpendicular to both ends and passing through its axis is a surface rectangular in shape. So it seems that the cylinder is both circular and rectangular at the same time. But I think it is clear what we should say here. The cylinder is not circular, nor is it rectangular. It is cylindrical. In fact, it doesn't even make sense to say of the *three-dimensional* cylinder that it is circular or rectangular, because 'circular' and 'rectangular' are predicates that only truly apply to *two-dimensional* objects. What *is* circular is a certain two-dimensional cross section of the cylinder, and what *is* rectangular is a distinct two-dimensional cross section. There is not one thing that is both circular and rectangular. There are at least three things that we are talking about here: a surface that is circular, another surface that is rectangular, and a solid that is cylindrical.

But can that be right? Isn't there only one thing, the cylinder, whose shape we are talking about? We say that in one sense *the cylinder* is circular and in another sense *the cylinder* is rectangular, and we seem not to think that there is anything else whose shape

we are talking about. I think the sensible thing to say is that there actually are things distinct from the cylinder, whose shapes we are describing, but they are *parts* of the cylinder. Just as my hand has many parts - fingers, nails, knuckles, and so on - the cylinder has many parts, which include all of its cross-sectional surfaces. And just as I can talk about my hand by talking about its parts (“My middle finger is longer than my index finger”), we can talk about the cylinder by talking about its parts (“This cross section is circular but that one is rectangular”).

I suggest we say a similar thing in the case of the way the cylinder *looks* in respect of shape. The way the cylinder looks from where Bill is and the way the cylinder looks from where Ben is are distinct from the way the cylinder looks, but the first two are parts of the last. Just as my hand has many parts, and the cylinder has many parts, the way the cylinder looks has many parts, which include the way it looks from where Bill is, the way it looks from where Ben is, and the way it looks from any other location in the room. And just as I can talk about my hand by talking about its parts, and about the cylinder by talking about its parts, we can talk about the way the cylinder looks by talking about its parts (“From here it looks circular but from there it looks rectangular”).

I will take it, then, that there are ways of looking such that the one object at the one time in the one possible world can have more than one of them, but such that the one object *from the one location* at the one time in the one possible world can have at most one of them. Then there is a function $Looks_1$ which takes these ways of looking as values, for which if x is an object, z is a location relative to x , t is a time, and Ω is a possible world then $Looks_1(x,z,t,\Omega)$ is the way that x looks from z at t in Ω .

But we seem to talk about ways of looking as though the one object from the one location at the one time in the one possible world can have more than one of them. The way a coloured object looks when viewed in the periphery of one’s visual field can be different from the way it looks when viewed in the centre. You and I are standing at the same location (near enough), looking at a coloured patch on a wall. I am looking at it out of the corner of my eye and report that it looks blue. You are looking at it out of the centre of your eyes and report that it looks red. We think that there is no inconsistency in what we have said. But if I were to say to you that out of the corner of my eye it looks both blue and red (all over), you would take me to be uttering an inconsistency. And if you were to say to me that out of the centre of your eyes it looks both red and

blue, I would take you to be uttering an inconsistency. We take it to be inconsistent to talk about a single way of looking by saying that its object both looks red and looks blue - to do so would be to say inconsistent things about that way of looking. But in this case we take there to be no disagreement between us. So we must take ourselves to be talking about distinct things: things that we refer to as ‘the way the patch looks out of the corner of my eye’ and ‘the way the patch looks out of the centre of my eye’.

I suggest that we say, again, that these are both distinct from, but parts of, the way the patch looks from where I am. I will take it, then, that there are ways of looking such that the one object from the one location at the one time in the one possible world can have more than one of them, but such that the one object from the one location *when viewed in the one way* at the one time in the one possible world can have at most one of them. Then there is a function $Looks_2$ which takes these ways of looking as values, for which if x is an object, z is a location relative to x , v is a way of viewing, t is a time, and Ω is a possible world then $Looks_1(x,z,v,t,\Omega)$ is the way that x looks from z when viewed in way v at t in Ω .

But we seem to talk about ways of looking as though the one object from the one location when viewed in the one way at the one time in the one possible world can have more than one of them. You and I, still standing at the same location, both now look from the centre of our eye at a Necker cube diagram on the wall. I report that it looks like a solid cube whose front face is lower than its back face, and you report that it looks like a solid cube whose front face is higher than its back face, and we are happy to accept that we have both said something true. Arguing similarly, we must take ourselves to be talking about distinct things, things we might refer to as ‘the way the diagram looks when apprehended downwards’ and ‘the way the diagram looks when apprehended upwards’.

I suggest that we say, again, that these are both distinct from, but parts of, the way the patch looks from where I am when viewed in a certain way. I will take it, then, that there are ways of looking such that the one object from the one location when viewed in the one way at the one time in the one possible world can have more than one of them, but such that the one object from the one location when viewed in the one way *and ‘apprehended’ in the one way* at the one time in the one possible world can have at most one of them. Then there is a function $Looks_3$ which takes these ways of looking as

values, for which if x is an object, z is a location relative to x , v is a way of viewing, a is a way of apprehending, t is a time, and Ω is a possible world then $Looks_3(x,z,v,a,t,\Omega)$ is the way that x looks from z when viewed in way v and apprehended in way a at t in Ω .

There may be more structure to ways of looking. That is, we may talk as if the one object from the one location when viewed in the one way and apprehended in the one way at the one time in the one possible world can have more than one way of looking. I will assume, however, that there are things I will call *aspects* such that there is at most one way that an object can look from the one aspect at the one time in the one possible world, and that these ways of looking have only respects as parts. From what I have said above we can say that aspects are multi-dimensional, with components ranging over relative locations, ways of viewing, and ways of apprehending: an aspect is an ordered n -tuple, for some n at least as great as three, whose components at least include one for a location, one for a way of viewing, and one for a way of apprehending. What is n , and what other components does an aspect include? As big as it takes, and whichever components it takes for object-aspect-time-world 4-tuples to uniquely determine a way of looking that does not itself have ways of looking as parts. From here on I will take ways of looking to be such that the one object from the one aspect at the one time in the one possible world can have at most one of them, and I will take the function *Looks* to be defined on object-aspect-time-world 4-tuples such that $Looks(x,a,t,\Omega)$ is the way that x looks from a at t in Ω .

Ways of looking, then, are complex things that have parts. Those parts can be ranked into levels, and at the lowest level we have ways of looking from a given aspect. I will take it that we can see ways of looking of all levels, but ultimately only in virtue of being able to see their lowest-level parts. And I will take it that if x is an object that looks some way to an observer S then there is a unique aspect a such that S can see the way that x looks from a , and ‘the way that x looks to S ’ refers to the way that x looks from aspect a . I shall sometimes refer to this aspect as *the aspect determined by S* . There may, for all we know, be ways of looking at levels lower than ways of looking from aspects. That is, contrary to what I have claimed above, ways of looking from aspects may themselves have (proper) parts. But if they do then we cannot see them, and that is why I propose we take ways of looking from aspects to be in the lowest level. Indeed,

aspects are intended to include just as many coordinates as are needed to determine a way of looking that we can see but whose parts, if it has any, we cannot.

Chapter 2

Talking About How Things Look

Frank Jackson (1977, ch. 2) distinguishes what he describes as three uses of the word ‘looks’ that he calls *epistemic*, *comparative*, and *phenomenal*. Examples of the epistemic use can be cast in the form ‘It looks as if p’ for some proposition p, and include ‘It looks as if these tomatoes are ripe’ and ‘It looks as if it is about to rain’. Examples of the comparative use can be cast in the form ‘X looks like an F normally does in C to S’ for some object X, kind of object F, circumstance C and person S, and include ‘It looks like a cow’ and ‘It looks like a red thing does at dusk’. Examples of the phenomenal use can be cast in the form ‘X looks F to S’ for some object X, colour/shape/distance predicate F, and person S, and include ‘It looks blue to me’ and ‘The top line looks longer than the bottom line’.⁸

I want to argue for a similar three-way distinction, and agree with Jackson in placing each of these pairs of examples into different classes based upon that distinction. But the distinction that I will make is not based upon different uses of the word ‘looks’. Rather, I will argue that ‘looks’ plays the same semantic role in all of these sentences - to express the function *Looks* and thereby to help determine a way of looking. The distinction I want to make is based upon the way in which these sentences are used to say something about that way of looking. There are, I will argue, three distinct ways in which they do so. If my analysis is right, then to say that in the sentences ‘It looks as if these tomatoes are ripe’, ‘It looks like a cow’, and ‘It looks blue to me’ we have three different uses of the word ‘looks’ is like saying that in the sentences ‘My dog is old’, ‘Your dog barks a lot’ and ‘His dog is as hairy as a goat’ we have three different uses of the word ‘dog’. I think a more accurate description of the matter is that in the first three sentences we have three different ways of talking about the way things look, just as in the last three sentences we have three different ways of talking about dogs. So that is how I shall describe it.

2.1 The Theory

We talk as if there are ways of looking. We ought to expect, then, that we talk as if describing and relating them. According to the theory that I want to develop, it is in our ‘looks’ talk that we do so. I will argue that, in particular, when we utter a sentence of the form ‘x looks ...’ for some object x, we are either describing the way that x looks or comparing it to some other way of looking, by expressing a proposition of one of the following three distinct forms:

$$(F1) \quad f_R(\text{Looks}(x,a,t,\Omega)) = r,$$

where x is an object, a is an aspect, t is a time, Ω is a world, R is some respect, r is an element of R, and f_R is a function from ways of looking to elements of R.

$$(F2) \quad f_Q(\text{Looks}(x,a,t,\Omega),x,a,t,\Omega) = p,$$

where Q is some question, p is an answer to Q, and f_Q is a function from ways of looking-object-aspect-time-world 5-tuples to propositions.

$$(F3) \quad f_R(\text{Looks}(x,a,t,\Omega)) \Phi f_R(w),$$

where R is some respect, Φ is some relation on the elements of R, f_R is a function from ways of looking to elements of R, and w is some way of looking. More generally:

$$(\forall R)[R \in \mathfrak{R} \rightarrow f_R(\text{Looks}(x,a,t,\Omega)) \Phi f_R(w)],$$

where \mathfrak{R} is some class of respects.

In the next three sections I will argue that in uttering sentences of the form ‘x looks ...’ we do indeed express propositions of these three forms.

2.2 One Way of Talking

I see a patch on the wall in front of me, and you ask me to say something about the way it looks. So I tell you that it looks red.

In so doing, I take myself to have used ‘red’ to refer to the *colour* red. I might have said, more fully, that the patch looks red in respect of colour rather than in respect of

⁸ These examples are all ones given by Jackson, pp. 30-3.

shape, because I take it that red is a colour and not a shape. If I had said that it looks square then I might have said, more fully, that it looks square in respect of shape rather than in respect of colour, because I take it that square is a shape and not a colour. If you were to ask me how the patch looks in respect of colour, or which colour the patch looks, I would take it that you want me to say which colour C it is such that the patch looks C . If I were to tell you which shape S it is such that the patch looks S , then I would not be answering your question. I might answer your question by simply saying 'red', taking myself to refer to the colour about which you are inquiring, or more fully by saying 'the patch looks red', using 'red' to do the same thing.

The same patch may look a different colour from a different aspect, or from the same aspect at a different time, or from the same aspect and at the same time but in a different possible world. But we take it that from a given aspect a , at a given time t , and in a given world Ω there can only be one colour C such that the patch looks C from a at t in Ω . If in addition to saying that the patch looks red (to me, now, in this world) I were to say that the patch also looks green, we would take my claims to be inconsistent, reasoning that the patch cannot look both red and green (all over). Well, almost. We would allow that the patch can look both red and some more specific shade of red, call it red_a . But this is because red *includes* red_a : the latter is a more specific colour than the former, in the sense that anything that is red_a must also be red, but something might be red without being red_a (it might be some other more specific shade of red). Red and red_a are not mutually exclusive. But some colours are. Red and green are mutually exclusive, because anything that is red cannot also be green, and vice versa. If C_1, \dots, C_n are pairwise mutually exclusive colours, then anything that is C_i for some i cannot also be C_j for any $j \neq i$. We take it that the same holds for the way things *look* in respect of colour: if C_1, \dots, C_n are pairwise mutually exclusive colours, then anything that *looks* C_i for some i cannot also *look* C_j for any $j \neq i$. If I were to say that the patch looks C_i and also looks C_j for distinct i and j , we would take my claims to be inconsistent. Moreover, you might ask me to talk about *the* way the patch looks in respect of colour, or to tell you *the* colour that it looks.

So we take it that there is only one colour C such that the patch looks C from a at t in Ω . But this is not because of anything special about the patch: if all I told you was that I can see an object, you might still ask me to tell you *the* colour that it looks, or

which colour it looks (note: colour, not colours), showing that we take it that *nothing* can look two (mutually exclusive) colours at the one time. We take it, then, that there is a function from object-aspect-time-world 4-tuples to colours, whose value for object x , aspect a , time t , and world Ω is the colour that x looks from a at t in Ω : that colour C such that x looks C in respect of colour from a at t in Ω . It may only be partially defined: there may be an x , a , t and Ω such that there is no colour C such that x looks C from a at t in Ω (is there a colour C such that the number 2 looks C from some aspect at some time in some world?). But it is, we take it, well-defined: for any given x , a , t , and Ω there is at most one colour C such that x looks C from a at t in Ω (on the understanding that we restrict ourselves to colours in some pairwise mutually exclusive class). If we call this function ' f ' then to say that x looks C in respect of colour is to say that $f(x,a,t,\Omega) = C$, for some aspect a (probably the aspect of the utterer), some t (probably the time of utterance), and some Ω (probably the world of utterance).

But we take it that there is more structure to what I have said than that. Having said that the patch looks red, I add that it also looks square. I could argue in a similar way to the conclusion that there is a function from object-aspect-time-world 4-tuples to shapes, whose value for the object x , aspect a , time t , and world Ω is the shape that x looks from a at t in Ω : that shape S such that x looks S in respect of shape from a at time t in Ω . If we call this function ' g ', then to say that x looks S in respect of shape is to say that $g(x,a,t,\Omega) = S$. So when I say that the patch looks red and that the patch looks square, I have said that $f(x,a,t,\Omega) = \text{red}$ and $g(x,a,t,\Omega) = \text{square}$. This explains why I take myself to have said something about *the patch* in each case – in each case the patch, x , is a constituent of what I have said. But it does not explain why I take myself to have said something about *the way the patch looks* in each case - there is no common constituent in $f(x,a,t,\Omega) = \text{red}$ and $g(x,a,t,\Omega) = \text{square}$ that is the way the patch looks. The only common constituents are x (the patch), a (an aspect), t (a time), and Ω (a world). So there must be more to what I have said than that $f(x,a,t,\Omega) = \text{red}$ and that $g(x,a,t,\Omega) = \text{square}$.

There is, we take it, a function, call it f_{Colour} , from ways of looking to colours, defined as follows: if w is a way of looking, and if there is an x , a , t , Ω and C such that w is the way that x looks from a at t in Ω and x looks C in respect of colour from a at t

in Ω , then $f_{Colour}(w) = C$; and if there is no such x, a, t, Ω and C then f_{Colour} is undefined on w . f_{Colour} may only be partially defined: there may be a way of looking for which there is no such x, a, t, Ω and C (if there is a way that the number 2 looks, then is f_{Colour} defined on that way?). But it is, we take it, well-defined. For suppose that w is a way of looking and that $f_{Colour}(w) = C$ for some C . Then there is an x, a, t , and Ω such that w is the way that x looks from a at t in Ω and x looks C in respect of colour from a at t in Ω . Since colour is, we take it, a respect of ways of looking, if x', a', t' , and Ω' are such that the way that x' looks from a' at t' in Ω' is identical to the way that x looks from a at t in Ω then x' also looks C from a' at t' in Ω' (more simply: if the way that x looks is identical to the way that x' looks then if x looks C in respect of colour then so does x'). So f_{Colour} is well-defined.

Similarly, there is, we take it, a function, call it f_{Shape} , from ways of looking to shapes, defined as follows: if w is a way of looking, and if there is an x, a, t, Ω and S such that w is the way that x looks from a at t in Ω and x looks S in respect of shape from a at t in Ω , then $f_{Shape}(w) = S$; and if there is no such x, a, t, Ω and S then f_{Shape} is undefined on w . f_{Shape} , also, may only be partially defined, but it is, we take it, well-defined.

The patch looks red. In virtue of what does it do so? The natural answer is: in virtue of the way that it looks. The patch looks a certain way, and it is in virtue of looking that way that it looks red. If another patch looked the (identically) same way, then it too would look red. If I were to say that this patch looks red, that the way that that patch looks is identical to the way that this patch looks, but that that patch does not look red, then we would take me to have made inconsistent claims.

I suggest, then, that the function f factors into the function *Looks*, defined above, and the function f_{Colour} : $f = f_{Colour} \circ Looks$ (the composition of f_{Colour} and *Looks*). And I suggest that the function g factors into the function *Looks* and the function f_{Shape} : $g = f_{Shape} \circ Looks$. Then to say that x looks C in respect of colour from a at t in Ω is to say that $f_{Colour}(Looks(x,a,t,\Omega)) = C$, and to say that x looks S in respect of shape from a at t in Ω is to say that $f_{Shape}(Looks(x,a,t,\Omega)) = S$. This explains why I take myself to have said something about *the patch* in each case: x , the patch, is a common constituent. It explains why I take myself to have said something about *the way the patch looks* in each

case: $Looks(x,a,t,\Omega)$, the way the patch looks from a at t in Ω , is a common constituent. It explains why I take myself to have said, in the first case, something about the way the patch looks in respect of *colour*, and in the second case something about the way the patch looks in respect of *shape*: the function f_{Colour} is a constituent of the first, and the function f_{Shape} is a constituent of the second. And, finally, it explains why I take it that if x looks red from a at t in Ω ($f_{Colour}(Looks(x,a,t,\Omega)) = \text{red}$) and if the way that y looks from a' at t' in Ω' is identical to the way that x looks from a at t in Ω ($Looks(y,a',t',\Omega') = Looks(x,a,t,\Omega)$) then y looks red from a' at t' in Ω' ($f_{Colour}(Looks(y,a',t',\Omega')) = \text{red}$), and why I take it that if x looks square from a at t in Ω ($f_{Shape}(Looks(x,a,t,\Omega)) = \text{square}$) and if the way that y looks from a' at t' in Ω' is identical to the way that x looks from a at t in Ω ($Looks(y,a',t',\Omega') = Looks(x,a,t,\Omega)$) then y looks square from a' at t' in Ω' ($f_{Shape}(Looks(y,a',t',\Omega')) = \text{square}$).

I have argued that ‘The patch looks red’ expresses a proposition $f_{Colour}(Looks(x,a,t,\Omega)) = \text{red}$, and that ‘The patch looks square’ expresses a proposition $f_{Shape}(Looks(x,a,t,\Omega)) = \text{square}$. The argument will extend to any sentence of the form ‘x looks F’ so long as (i) ‘x’ is a singular term, (ii) ‘F’ refers to an element of some respect R, and (iii) if x looks F and if the way that y looks is identical to the way that x looks then it follows that y looks F, a condition that will hold if R is any respect of ways of looking (so that ‘looks F’ is an observational predicate). If these conditions are met then ‘x looks F’ expresses a proposition of the form:

$$(F1) \quad f_R(Looks(x,a,t,\Omega)) = F,$$

where f_R is a function from ways of looking to elements of R.

There may well be more structure to the functions f_{Colour} , f_{Shape} , and, more generally, f_R . There may, for instance, be a function f such that $f_{Colour} = f(\text{Colour})$, $f_{Shape} = f(\text{Shape})$, and more generally $f_R = f(R)$. A contender for such an f might be: that function such that if R is a respect of ways of looking and if w is a way of looking then $f(R)(w)$ is that element r of R such that most x for which x is r are such that x has w (so that to say that x looks r in respect of R is to say that x looks the way most r things look).⁹ If there is

⁹ I don’t think that such a statistical function is a *plausible* contender. Suppose that to say that x looks red is to say that x looks the way most red things look. Then if x looks red in the actual world and W is a possible world in which x looks the same as it does in the actual world but in which most red things look blue, then it would be false that x looks red in W. But that is not what we think – we think that it would be true that x looks red in W. The example is merely illustrative.

such a function, then by grasping it and by grasping the respects of colour, shape, and so on, we can thereby grasp the functions f_{Colour} , f_{Shape} , and so on: anyone who grasps the function f and the respect of, say, length can thereby grasp the function f_{Length} . The extent to which this is true is the extent to which we should think that there is such further structure. I will not pursue the matter here. For present purposes, all I need is the result that some sentences of the form ‘ x looks ...’ express propositions of the form F1, whether or not those propositions are also instances of a more structured form.

2.3 A Second Way of Talking

You ask me to say something more about the way the patch looks. I have already told you that it looks red. I believe that lighting conditions are such that only red patches look red (suppose that there’s a white light source illuminating the wall). So now I tell you that it looks as though it *is* red.

I cannot thereby have expressed the same proposition, $f_{Colour}(Looks(x,a,t,\Omega)) = \text{red}$, that I did when I said that the patch looks red. For had I believed that lighting conditions are such that only white patches look red, then I would still have told you that the patch looks red, but I would have denied that it looks as though it is red, and told you instead that it looks as though it is white. Although the propositions that the patch looks red and that the patch looks as though it is red have the same truth value in the actual world, they have different truth values in the counterfactual world, so must be distinct propositions.

Nor can I have expressed a proposition $f_{Colour}(Looks(x,a,t,\Omega)) = C$ for some other colour C . For that would be to say that the patch looks C . If that is right, then in the counterfactual world above I would still have believed that the patch looks C , and hence ought still have told you that the patch looks as though it is red. But I wouldn’t have – I would have told you that the patch looks as though it is white. More generally, I cannot have expressed a proposition $f_R(Looks(x,a,t,\Omega)) = r$ for *any* respect R . For that would be to say that the patch looks r in respect of R , in which case I ought still have told you in the counterfactual world that the patch looks as though it is red. But, again, I wouldn’t have.

So when I say that the patch looks as though it is red, I am expressing a proposition that is not of the form F1. We have here a second way of talking about the way things look.

So what *have* I said? I take it that I have used the words 'it is red' to express the proposition that the patch is red. Having said that the patch looks as though it is red, I might have added, 'but that's not the case'. What's not the case? That the patch is red. When I say 'but that's not the case' I intend the word 'that' to refer to the proposition that I expressed with the words 'it is red'. So I take those words to have expressed a proposition and I take them to have expressed the proposition that the patch is red.

The patch looks red. It may also look, from the same aspect at the same time in the same world, square and large. There may be many Fs such that the patch looks F. But in respect of *colour* there is only one colour C such that the patch looks C (assuming that colours are pairwise mutually exclusive). I think we can make similar remarks in this case as well. The patch looks as though it is red. It may also look, from the same aspect at the same time in the same world, as though it is square, as though it has spots on it, as though someone took care in making it, and as though the wall is being illuminated by white light. There may be many propositions p such that the patch looks as though p. But in respect of *which colour it is* there is, I suggest, only one proposition p such that the patch looks as though p. I have said that in respect of which colour it is the patch looks as though *it is red*. If I had added that in respect of which colour it is the patch also looks as though *it is green* we would take me to have made inconsistent claims. In this case I suggest that we take the respect to be a *question*: 'What colour is the patch?'. And I suggest that we take the elements of this respect to be possible *answers* to this question: 'The patch is red', 'The patch is yellow', and so on. By saying that the patch looks as though it is red I have singled out an answer to this question.

We take it that there is only one proposition, p, such that in respect of which colour it is the patch looks as though p. But this is not because of anything special about the patch: if all I told you was that I can see an object, you might still ask me to tell you *the* way that it looks in that respect, or which way it looks (note: way, not ways), showing that we take it that *nothing* can look two ways in respect of which colour it is at the one time. We take it, then, that there is a function from object-aspect-time-world 4-tuples to propositions, whose value for object x, aspect a, time t and world Ω is the proposition p

such that in respect of which colour it is x looks as though p from a at t in Ω . It may only be partially defined: there may be an x , a , t and Ω for which there is no such proposition. But it is, we take it, well-defined: for any given x , a , t and Ω there is at most one proposition p such that in respect of which colour it is x looks as though p from a at t in Ω . If we call this function ‘ C ’, then to say that in respect of which colour it is x looks as though p is to say that $C(x,a,t,\Omega) = p$, for some a (probably the aspect of the utterer), some t (probably the time of utterance) and some Ω (probably the world of utterance).

But we take it that there is more structure to what I have said than that. Having said that the patch looks as though it is red, I add that it also looks as though it is square. I could argue in a similar way to the conclusion that there is a function from object-aspect-time-world 4-tuples to propositions, whose value for the object x , aspect a , time t and world Ω is the proposition p such that in respect of which shape it is x looks as though p from a at t in Ω . If we call this function ‘ S ’, then to say that in respect of which shape it is x looks as though p is to say that $S(x,a,t,\Omega) = p$. So when I say that the patch looks as though it is red and that the patch looks as though it is square, I have said that $C(x,a,t,\Omega) = p_c$ and $S(x,a,t,\Omega) = p_s$, where p_c is the proposition that the patch is red and p_s is the proposition that the patch is square. This explains why I take myself to have said something about *the patch* in each case – in each case the patch, x , is a constituent of what I have said. But it does not explain why I take myself to have said something about *the way the patch looks* in each case - there is no common constituent of $C(x,a,t,\Omega) = p_c$ and $S(x,a,t,\Omega) = p_s$ that is the way the patch looks. The only common constituents are x (the patch), a (an aspect), t (a time), and Ω (a world). So there must be more to what I have said than that $C(x,a,t,\Omega) = p_c$ and that $S(x,a,t,\Omega) = p_s$.

In the previous section I argued that the functions f and g have a common factor, in virtue of there being the function f_{Colour} from ways of looking to colours and the function f_{Shape} from ways of looking to shapes. To get a common factor in a similar way here, we need there to be a function, call it f_{Qc} (‘ Qc ’ for ‘question about colour’), from ways of looking to propositions, defined as follows: if w is a way of looking, then if there is an object x , aspect a , time t , world Ω , and proposition p such that w is the way that x looks from a at t in Ω and in respect of which colour it is x looks as though p from

a at t in Ω , then $f_{Q_c}(w) = p$; if there is no such x, a, t, Ω and p then f_{Q_c} is undefined on w. But there is no such function, for f_{Q_c} is not well-defined: in respect of which colour it is, the patch x looks as though it is red (to me, now, in this world), but the patch y looks as though it is white, even though the way that x looks is identical to the way that y looks. Nor is there an appropriate function from ways of looking–object pairs to propositions, because the one object with the one way of looking may at one time look as though it is red but at another look as though it is white (if, for example, the lighting conditions were different at those times). Nor is there an appropriate function from ways of looking–object–aspect triples to propositions, because the one object with the one way of looking from the one aspect may at one time look as though it is red but in another look as though it is white (if, for example, the lighting conditions were adjusted accordingly). Nor is there an appropriate function from ways of looking–object–aspect–time 4-tuples to propositions, because the one object with the one way of looking from the one aspect at the one time may in one world look as though it is red but in another look as though it is white (if, for example, the lighting conditions were different in those worlds).

But there *is* an appropriate function from ways of looking–object–aspect–time–world 5-tuples to propositions, because for the one object with the one way of looking from the one aspect at the one time in the one world there is only one proposition p such that in respect of which colour it is that object looks as though p from that aspect at that time in that world. So f_{Q_c} is well-defined if we define it as follows: if w is a way of looking, x is an object, a is an aspect, t is a time and Ω is a world, then if w is the way that x looks from a at t in Ω and if there is a proposition p such that in respect of which colour it is x looks as though p from a at t in Ω , then $f_{Q_c}(w,x,a,t,\Omega) = p$; otherwise f_{Q_c} is undefined on (w,x,a,t,Ω) . f_{Q_c} may only be partially defined: w may not be the way that x looks from a at t in Ω , and there may be no such p. But it is, we take it, well defined. Similarly, there *is* a function, call it f_{Q_s} ('Q_s' for 'question about shape'), from ways of looking–object–aspect–time–world 5-tuples to propositions, defined as follows: if w is a way of looking, x is an object, a is an aspect, t is a time and Ω is a world, then if w is the way that x looks from a at t in Ω and if there is a proposition p such that in respect of which shape it is x looks as though p from a at t in Ω , then $f_{Q_s}(w,x,a,t,\Omega) = p$; otherwise

f_{Q_s} is undefined on (w,x,a,t,Ω) . Again, f_{Q_s} may only be partially defined, but it is, we take it, well defined.

I suggest, then, that what I am saying when I say that the patch looks as though it is red is not that $C(x,a,t,\Omega) = p_c$, but that $f_{Q_c}(Looks(x,a,t,\Omega),x,a,t,\Omega) = p_c$ (where a is the aspect of utterance, t is the time of utterance and Ω is the world of utterance); and that what I am saying when I say that the patch looks as though it is square is not that $S(x,a,t,\Omega) = s$, but that $f_{Q_s}(Looks(x,a,t,\Omega),x,a,t,\Omega) = p_s$. This explains why I take myself to have said something about *the patch* in each case: x , the patch, is a common constituent. It explains why I take myself to have said something about *the way the patch looks* in each case: $Looks(x,a,t,\Omega)$, the way the patch looks, is a common constituent. It explains why I take myself to have said, in the first case, something about the way the patch looks in respect of *which colour it is*, and in the second case something about the way the patch looks in respect of *which shape it is*: the first has the function f_{Q_c} as a constituent, and the second has the function f_{Q_s} as a constituent. And, finally, it does not license the inferences that we think ought not be licensed: if x looks as though it is red ($f_{Q_c}(Looks(x,a,t,\Omega),x,a,t,\Omega) = p_c$) and if the way that y looks is identical to the way that x looks ($Looks(y,a,t,\Omega) = Looks(x,a,t,\Omega)$) then it does not follow that y looks as though it is red ($f_{Q_c}(Looks(y,a,t,\Omega),y,a,t,\Omega) = p_c$), and if x looks as though it is square ($f_{Q_s}(Looks(x,a,t,\Omega),x,a,t,\Omega) = p_s$) and if the way that y looks is identical to the way that x looks ($Looks(y,a,t,\Omega) = Looks(x,a,t,\Omega)$) then it does not follow that y looks as though it is square ($f_{Q_s}(Looks(y,a,t,\Omega),y,a,t,\Omega) = p_s$).

I have argued that ‘The patch looks as though it is red’ expresses a proposition $f_{Q_c}(Looks(x,a,t,\Omega),x,a,t,\Omega) = p_c$, and that ‘The patch looks as though it is square’ expresses a proposition $f_{Q_s}(Looks(x,a,t,\Omega),x,a,t,\Omega) = p_s$. The argument will extend to any sentence of the form ‘ x looks as though p ’ so long as (i) ‘ x ’ is a singular term, and (ii) ‘ p ’ expresses a proposition. If these conditions are met then ‘ x looks as though p ’ expresses a proposition of the form:

$$(F2) \quad f_Q(Looks(x,a,t,\Omega),x,a,t,\Omega) = p,$$

where Q is some question, p is an answer to Q , and f_Q is a function from ways of looking-object-aspect-time-world 5-tuples to propositions.¹⁰

I will take it that sentences of the form ‘ x looks as if p ’ and ‘ x looks like p ’ can be cast in the form ‘ x looks as though p ’, and hence express propositions of the form F2. Moreover, I will take it that sentences of the form ‘ x looks to be F ’ can be cast in the form ‘ x looks as though it is F ’, that sentences of the form ‘ x looks to have y ’ can be cast in the form ‘ x looks as though it has y ’, that sentences of the form ‘ x looks to want y ’ can be cast in the form ‘ x looks as though it wants y ’, and so on, and that these all express propositions of the form F2 as well.

There may well be more structure to the functions f_Q . There may, for instance, be a function f such that $f_Q(Looks(x,a,t,\Omega),x,a,t,\Omega) = f(Q,Looks(x,a,t,\Omega),x,a,t,\Omega)$. A contender for such an f might be: that function such that if Q is a question, x is an object, a is an aspect, t is a time, and Ω is a world, then $f(Q,Looks(x,a,t,\Omega),x,a,t,\Omega)$ is that answer p to the question Q such that the fact that the way that x looks from a at t in Ω is $Looks(x,a,t,\Omega)$ is evidence that p .¹¹ If there is such a function, then by grasping it we can thereby grasp the functions f_Q for any question Q : if Q_1 is the question ‘What colour is the patch’ and Q_2 is any other question, then anyone who grasps the function f_{Q_1} can thereby grasp the function f_{Q_2} . The extent to which this is true is the extent to which we should think that there is such further structure. I will not pursue the matter here. For present purposes all I need is the result that some sentences of the form ‘ x looks ...’ express propositions of the form F2 and not propositions of the form F1, whether or not those propositions are also instances of a more structured form.

2.4 A Third Way of Talking

You ask me to say something more about the way the patch x looks. I have already told you that it looks red and that it looks as though it is red. I can also see another patch

¹⁰ Technically, the first argument of the function f_Q is redundant: its value is determined by the other four. I am not happy with this. My aim has been to find structurally similar forms for these two classes of propositions, but this redundancy suggests that I have not yet succeeded. More work needs to be done.

¹¹ Again, I am not claiming that this function is a *plausible* contender. It is not clear, for example, that it is even well defined – that for a given question Q it gives a unique answer p . More needs to be said, also, about the notion of ‘evidence’ being employed here. The example is merely illustrative.

y on a second wall such that y also looks red. So now I tell you that x looks the same as y in respect of colour.

I cannot thereby have expressed a proposition $f_R(Looks(x,a,t,\Omega)) = r$ for some respect R and some member r of R. For then had I seen that y looks yellow rather than red I ought still to have told you that x looks the same as y in respect of colour. But I would not have. So to say that x looks the same as y in respect of colour cannot be to express a proposition of the form F1.

Nor can I have expressed a proposition $f_Q(Looks(x,a,t,\Omega),x,a,t,\Omega) = p$ for some question Q and answer p. For if that is what I have said then it does not follow that if z is an object such that the way that z looks is identical to the way that x looks then z looks the same as y in respect of colour. But that is something that we *do* take to follow from what I have said. If I were to say that x looks the same as y in respect of colour, that the way that z looks is identical to the way that x looks, but that z does not look the same as y in respect of colour, then we would take me to have made inconsistent claims. So to say that x looks the same as y in respect of colour cannot be to express a proposition of the form F2.

So when I say that x looks the same as y in respect of colour, I am not expressing a proposition of the form F1 or F2. We have here a third way of talking about the way things look.

So what *have* I said? I said that x looks the same as y in respect of colour. I take myself to have said something about x (that it looks the same as y in respect of colour), something about y (that x looks the same as it in respect of colour), and something about the way that x looks. But I also take myself to have said something about the way that y looks: First, I might have said more fully that x looks the same as y *looks* in respect of colour. Second, I take it that I could have said the same thing by saying that y looks the same as x in respect of colour, making it clearer that I have said something about the way that y looks. So I take myself to have said something about the way that x and y look. If I were to also say that x *is* the same as y in respect of colour, I would take myself to have said the same thing about the way that x and y *are* as I did about the way that x and y *look*. I take it that in this second statement I have said that the colour that x *is* is the same as the colour that y *is* – that $f_{Colour}(Is(x,t,\Omega)) = f_{Colour}(Is(y,t,\Omega))$. If this is to say the same thing about the way that x and y *are* as I said in the first statement about

the way that x and y *look*, then in that statement I must have said that the colour that x *looks* is the same as the colour that y *looks* – that $f_{Colour}(Looks(x,a,t,\Omega)) = f_{Colour}(Looks(y,a,t,\Omega))$. So this, I suggest, is what I *did* say. It explains why I take myself to have said something about x , y , the way that x looks, and the way that y looks – x , y , $Looks(x,a,t,\Omega)$, and $Looks(y,a,t,\Omega)$ are all constituents of what I have said. It explains why I have said something about the way that x and y look in respect of *colour* (rather than, say, *shape*) – the function f_{Colour} (rather than the function f_{Shape}) is a constituent of what I have said. And it explains why I take myself to have said the same thing about the way that x and y *look* as I have about the way that x and y *are* when I say that x *is* the same as y in respect of colour – the propositions $f_{Colour}(Looks(x,a,t,\Omega)) = f_{Colour}(Looks(y,a,t,\Omega))$ and $f_{Colour}(Is(x,t,\Omega)) = f_{Colour}(Is(y,t,\Omega))$ have the same structure and constituency apart from the occurrence of ‘*Looks*’ in the first and ‘*Is*’ in the second.

The argument can be modified to work for other things that I might say about the way that x looks, where x is any object. I say that x looks *like* y in respect of colour. I take myself to have said the same thing about the way that x and y *look* as I would say about the way that x and y *are* if I were to say that x *is* like y in respect of colour. If the latter is to say that $f_{Colour}(Is(x,t,\Omega)) \approx f_{Colour}(Is(y,t,\Omega))$ (i.e. that the colours $f_{Colour}(Is(x,t,\Omega))$ and $f_{Colour}(Is(y,t,\Omega))$ are in the same similarity class), then the former is to say that $f_{Colour}(Looks(x,a,t,\Omega)) \approx f_{Colour}(Looks(y,a,t,\Omega))$. I say that x looks *longer than* y . I take myself to have said the same thing about the way that x and y *look* as I would say about the way that x and y *are* if I were to say that x *is* longer than y . If the latter is to say that $f_{Length}(Is(x,a,t,\Omega)) > f_{Length}(Is(y,a,t,\Omega))$, then the former is to say that $f_{Length}(Looks(x,a,t,\Omega)) > f_{Length}(Looks(y,a,t,\Omega))$. I say that x looks *rounder than* y . I take myself to have said the same thing about the way that x and y *look* as I would say about the way that x and y *are* if I were to say that x *is* rounder than y . If the latter is to say that $\delta(f_{Shape}(Is(x,a,t,\Omega)), Circle) < \delta(f_{Shape}(Is(y,a,t,\Omega)), Circle)$, where δ is some metric on shapes, then the former is to say that $\delta(f_{Shape}(Looks(x,a,t,\Omega)), Circle) < \delta(f_{Shape}(Looks(y,a,t,\Omega)), Circle)$. In each of these cases – ‘ x looks the same as y in respect of colour’, ‘ x looks like y in respect of colour’, ‘ x looks longer than y ’, ‘ x looks rounder than y ’ – there is a respect R and a relation Φ on the elements of R such that what I have said is that $f_R(Looks(x,a,t,\Omega)) \Phi f_R(Looks(y,a,t,\Omega))$. The respects are, respectively:

colour, colour, length and shape. The relations are, respectively: being identical to, being similar to, being greater than, and being closer to circular than.

When I say that x looks the same as y , I might not mean that it looks the same as y in a single respect (say, colour). I might mean that it looks the same in several respects, or that it looks the same in all respects of ways of looking (in which case I might clarify by saying that it looks exactly the same as y , or that it looks identical to y , or that x and y look identical, or that the way that x looks is identical to the way that y looks). I suggest that in these cases there is a class of respects \mathfrak{R} such that what I am saying is that $(\forall R)(R \in \mathfrak{R} \rightarrow f_R(Looks(x,a,t,\Omega)) = f_R(Looks(y,a,t,\Omega)))$ (in the case that x looks identical to y , \mathfrak{R} is the class of respects of ways of looking). Similarly, when I say that x looks like y , I might not mean that it looks like y in a single respect, but that it looks like y in several respects. I suggest that in this case there is a class of respects \mathfrak{R} such that what I am saying is that $(\forall R)(R \in \mathfrak{R} \rightarrow f_R(Looks(x,a,t,\Omega)) \approx f_R(Looks(y,a,t,\Omega)))$. What if I were to say that x looks like a dog (looks)? I suggest that there is a certain way of looking w that I have in mind, perhaps the way that most dogs look, or that some dogs look, or that some paradigm dog looks, and a class of respects \mathfrak{R} such that what I am saying is that $(\forall R)(R \in \mathfrak{R} \rightarrow f_R(Looks(x,a,t,\Omega)) \approx f_R(w))$.

I claim, then, that when we talk about the way that x looks in this third way we are expressing a proposition of the form:

$$(F3) \quad f_R(Looks(x,a,t,\Omega)) \Phi f_R(w),$$

where R is some respect, Φ is some relation on the elements of R , and w is some way of looking. More generally, we are expressing a proposition of the form:

$$(F3a) \quad (\forall R)(R \in \mathfrak{R} \rightarrow f_R(Looks(x,a,t,\Omega)) \Phi f_R(w)),$$

where \mathfrak{R} is some class of respects.

2.5 The Three Ways

There are, then, at least three distinct forms of proposition that we express using sentences of the form ‘ x looks ...’. There may be more, but much of what we say about the way things look we say, I think, by expressing propositions of these three forms.

But many of our ‘looks’ sentences are ambiguous, and some care needs to be taken in deciding exactly which propositions they express. I say that x looks red. I probably mean to express a proposition of the form F1, but I might mean to express a proposition of the form F2. Which one I mean should show up in how I’m prepared to say it more fully. If I’m prepared to say more fully that x looks *to be* red, or that x looks *like it is* red, or that x looks *as if it is* red, or that x looks *as though it is* red, then I probably mean to express a proposition of the form F2. If not, then I probably mean to express a proposition of the form F1. It should also show up in which inferences I take to be licensed by what I have said: if I take it that if y looks identical to x (at least in respect of colour) then it follows that y also looks red (when understood in the same sense), then I must mean to express a proposition of the form F1 (a proposition of that form licenses such an inference, but a proposition of the form F2 does not). These two tests can be used to disambiguate similar sentences: ‘x looks expensive’, ‘x looks drunk’, ‘x looks clean’, and so on.

I say that x looks the same as y. I probably mean to express a proposition of the form F3, but I might mean to express a proposition of the form F2. Again, which one I mean should show up in how I’m prepared to say it more fully. If I’m prepared to say more fully that x looks the same as y *looks*, or that x looks the same *way* as y, then I probably mean to express a proposition of the form F3. If I’m prepared to say more fully that x looks *to be* the same as y, or that x looks *like it is* the same as y, or that x looks *as if it is* the same as y, or that x looks *as though it is* the same as y, then I probably mean to express a proposition of the form F2. And it should also show up in which inferences I take to be licensed by what I have said: if I take it that if the way that z looks is identical to the way that x looks then it follows that z also looks the same as y, or that if the way that z looks is identical to the way that y looks then it follows that x also looks the same as z (when I understand ‘looks the same as’ in a constant sense), then I must mean to express a proposition of the form F3 (a proposition of that form licenses such inferences, but a proposition of the form F2 does not). These two tests can be used to disambiguate similar sentences: ‘x looks like y’, ‘x looks longer than y’, ‘x looks as tall as y’, and so on.

The class of sentences of the form ‘x looks ...’ that express a proposition of the form F1 is close to the class of sentences that Jackson calls *phenomenal* uses of ‘looks’.

The two do not exactly coincide, because the latter includes sentences like ‘The top line looks longer than the bottom line’ whereas the former does not (according to the theory developed here, such sentences express propositions of the form F3). Nevertheless, I am happy to call sentences that express a proposition of the form F1 *phenomenal* ‘looks’ sentences.

The class of sentences of the form ‘x looks ...’ that express a proposition of the form F2 is close to the class of sentences that Jackson calls *epistemic* uses of ‘looks’. More accurately, it is close to that subclass of the latter that have the form ‘x looks as if p’ rather than ‘it looks as if p’, because the theory developed here has, as yet, nothing to say about sentences that begin with a non-referential ‘it’. I am happy to call sentences that express a proposition of the form F2 *epistemic* ‘looks’ sentences.

The class of sentences of the form ‘x looks ...’ that express a proposition of the form F3 is close to the class of sentences that Jackson calls *comparative* uses of ‘looks’. Sentences in the latter class can be cast in the form ‘X looks like an F normally does in C to S’, and each of these is taken by the theory developed here to express a relation between the way that x looks and some other way of looking, and thus to express a proposition of the form F3. But the two classes do not exactly coincide, because the former includes sentences like ‘The top line looks longer than the bottom line’ whereas the latter does not (according to Jackson, they are *phenomenal* uses of ‘looks’). Nevertheless, I am happy to call sentences that express a proposition of the form F3 *comparative* ‘looks’ sentences.

2.6 Two Results

In this section I derive two results that are needed for Chapter 3.

2.6.1 The Transitivity of ‘Looks the Same As’

If x looks the same as y (to me, now, in this world) and y looks the same as z, does it follow that x looks the same as z?

That depends on what is being said. It *does*, if ‘x looks the same as y’ is saying that the way that x looks is identical to the way that y looks, and ‘y looks the same as z’ is

saying that the way that y looks is identical to the way that z looks, and ‘x looks the same as z’ is saying that the way that x looks is identical to the way that z looks. More precisely, it *does*, if for some aspects a, a', and a'', ‘x looks the same as y’ is saying that $Looks(x,a,t,\Omega) = Looks(y,a',t,\Omega)$, and ‘y looks the same as z’ is saying that $Looks(y,a',t,\Omega) = Looks(z,a'',t,\Omega)$, and ‘x looks the same as z’ is saying that $Looks(x,a,t,\Omega) = Looks(z,a'',t,\Omega)$. This follows from the transitivity of identity.

It *does*, if, for some respect R, ‘x looks the same as y’ is saying that the way that x looks in respect of R is identical to the way that y looks in respect of R, and ‘y looks the same as z’ is saying that the way that y looks in respect of R is identical to the way that z looks in respect of R, and ‘x looks the same as z’ is saying that the way that x looks in respect of R is identical to the way that z looks in respect of R. This follows, again, from the transitivity of identity.

It does *not*, if ‘x looks the same as y’ is saying that x looks as though it is y, and ‘y looks the same as z’ is saying that y looks as though it is z, and ‘x looks the same as z’ is saying that x looks as though it is z. For suppose that I can see of x that he looks the way that y normally does, that I can see of y that he looks the way that z normally does, and I believe that the way that y normally looks is quite different from the way that z normally looks. Then I might say that x (to me, now, in this world) looks as though he is y, that y looks as though he is z, but that x does *not* look as though he is z.

It does *not*, if, for some respect R, ‘x looks the same as y’ is saying that x looks as though it is the same as y in respect of R, and ‘y looks the same as z’ is saying that y looks as though it is the same as z in respect of R, and ‘x looks the same as z’ is saying that x looks as though it is the same as z in respect of R. For suppose that I can see distinct colour patches x, y and z and know (i) that x, y and z are under lighting conditions in which things usually look the colour that they are (i.e. for most x, if x *looks* C in respect of colour then x *is* C in respect of colour), (ii) that x looks red, that y is red but looks green, and that z is green. Then I might say that x looks as though it is the same as y in respect of colour, that y looks as though it is the same as z in respect of colour, but that x does *not* look as though it is the same as z in respect of colour.

The result that I need is this: ‘x looks the same as y’ can be used to express both a transitive and a non-transitive relation between x and y, so care needs to be taken when talking about the transitivity of ‘looks the same as’.

2.6.2 A Truism about ‘Looks the Same As’ and ‘Looks F’?

If x looks the same as y (to me, now, in this world) and x looks F, does it follow that y looks F?

That depends, again, on what is being said. It *does*, if ‘x looks the same as y’ is saying that the way that x looks is identical to the way that y looks, and, for some respect R and element r of R, ‘x looks F’ is saying that x looks r in respect of R and ‘y looks F’ is saying that y looks r in respect of R. More precisely, it *does*, if for some aspects a and a’, ‘x looks the same as y’ is saying that $Looks(x,a,t,\Omega) = Looks(y,a',t,\Omega)$, and ‘x looks F’ is saying that $f_R(Looks(x,a,t,\Omega)) = r$, and ‘y looks F’ is saying that $f_R(Looks(y,a',t,\Omega)) = r$. This follows from the well-definedness of the function f_R .

It *does*, if, for some respect R, ‘x looks the same as y’ is saying that the way that x looks in respect of R is identical to the way that y looks in respect of R, and, for some element r of R, ‘x looks F’ is saying that x looks r in respect of R, and ‘y looks F’ is saying that y looks r in respect of R. This follows from the transitivity of identity.

It does *not*, if, ‘x looks the same as y’ is saying that the way that x looks is identical to the way that y looks, and, for some respect R and element r of R, ‘x looks F’ is saying that x looks as though it is r in respect of R, and ‘y looks F’ is saying that y looks as though it is r in respect of R. For suppose that I can see distinct colour patches x and y and know (i) that x looks red and is in lighting conditions under which only red things look red, and (ii) that y looks red but is in lighting conditions under which only white things look red. Then I might say that the way that x looks is identical to the way that y looks, that x looks as though it is red, but that y does *not* look as though it is red (it looks as though it is white).

It does *not*, if, for some respect R, ‘x looks the same as y’ is saying that x looks as though it is the same as y in respect of R, and, for some element r of R, ‘x looks F’ is saying that x looks r in respect of R, and ‘y looks F’ is saying that y looks r in respect of R. For suppose that I can see distinct colour patches x and y and know (i) that x looks red but is in lighting conditions under which only white things look red, and (ii) that y looks green but is white. Then I might say that x looks as though it is the same as y in

respect of colour (i.e. white), that x looks red, but that y does not look red (it looks green).

The result that I need is this: 'x looks the same as y', 'x looks F' and 'y looks F' can be used to express propositions such that the truth of the first two does *not* entail the truth of the third, and to express propositions such that the truth of the first two *does* entail the truth of the third. So care needs to be taken when talking about whether or not this sentence expresses a truism: 'If x looks the same as y and x looks F, then y looks F.'

2.7 Visual Experience and its Content

There is currently some interest in perceptual experience and whether or not its content is conceptual or nonconceptual. I want to finish this chapter by considering some of what we should say about these matters, if we accept the theory that I have been developing. Because I have limited my attention to the way we talk about how things look rather than sound, taste, feel and smell, I will limit my attention to visual experience rather than perceptual experience via other sense modalities. And because I have only considered sentences of the form 'x looks ...' where 'x' is a singular referring expression, and not sentences of the form 'it looks as if ...' where 'it' is non-referential, I will only consider visual experiences that are experiences *of* objects (be they veridical or non-veridical) and will not consider visual experiences that are hallucinations.

It is standardly thought that when a subject S has a visual experience of an object x the experience represents x to be a certain way, and that to talk about the way that x looks to S is to talk about the way the experience represents x to be: to say that x looks red to S, for example, is to say that S's experience represents that x is red. This suggests that the content of the experience is propositional, which in turn suggests that it is conceptual - that the range of concepts that S possesses constrains the range of her possible visual experiences, just as it might be thought to constrain the range of her possible beliefs, desires and other propositional attitudes. Anyone who wants to argue that the content of visual experience is nonconceptual is better placed if he can put forward an account according to which it is not propositional.

Christopher Peacocke (1993) has done just that. He proposes that the content of a visual experience is a *scenario*: a set of ways of filling out the space around the subject of the experience that are consistent with the content's being correct (the experience being veridical). The content of the experience is correct just in case the space around the subject is occupied in one of those ways.¹² He claims that this is a suitable basis for a theory of nonconceptual content because a scenario is not 'built up' from concepts, and there is no requirement that its specification be limited to concepts possessed by the subject of the experience.¹³

According to the theory developed here, visual experience does not have content of this kind. To talk about the way that *x* looks to *S* is to talk about *a way of looking* – the way that *x* looks from *a*, where *a* is the aspect determined by *S*. If we accept that the content of *S*'s visual experience is whatever it is that we talk about when we talk about the way that *x* looks to *S*, it follows that the content of the experience is a way of looking. The experience does *not* represent that *x* is a certain way (or one of a set of ways), it represents the way that *x* looks.¹⁴ So to specify the content of the experience is *not* to specify the way that it represents *x* to be, because it does not represent *x* to be any way. Rather, it is to specify a way of looking. It might be thought that if *S*'s experience represents the way of looking that *x* has, then in virtue of that it does after all represent the way that *x* is, because that way of looking represents the way that *x* is. But that would be a mistake. Ways of looking are not representations. They are not things that we create or endow with representational properties. They are things that exist independently of us. And even though when an object has a certain way of looking it might be a *sign* that the object is a certain way, the way of looking does not *represent* that the object is that way, just as even though puddles of water on the ground might be a sign that there are clouds in the sky, the puddles do not represent that there are clouds.

If this is right, then why is it tempting to think that visual experience represents things to be a certain way? It is, I suggest, because of the particular way in which we talk about ways of looking. To illustrate, consider a possible way of talking about the ground. Suppose that we had no expressions to stand for properties that the ground can

¹² On some views, for the content of an experience to be propositional just is for it to have a correctness condition, and Peacocke's scenarios would thereby count as being propositional. I take it that Peacocke would reject such views.

¹³ See especially p. 63.

¹⁴ Note: not *which* way it looks, for that would be to represent a proposition rather than a way of looking.

have, so that we could not attribute properties to the ground using sentences of the form ‘the ground is ...’. Rather, seeing that we already have expressions to stand for all the properties that the sky can have we decide to use those instead. So with each property that the ground can have we associate a property that the sky can have. We notice that when the ground has puddles on it this is a sign that the sky is cloudy, so it is natural, but not mandatory, to associate the property of being puddled with the property of being cloudy. Then we can use that property to say that the ground is puddled. We cannot do so by saying that the ground is cloudy, because the property of being cloudy is not one that the ground can have. We cannot do so by saying that the sky is cloudy because the ground may be puddled even when the sky is not cloudy. But we can do so by saying that *the sky grounds cloudy*. Even though the form of words might suggest it, we would *not* be saying that the sky is represented as being cloudy, because the puddled ground does not represent that the sky is cloudy, even though it may be a sign that it is. We would be saying that the ground is puddled.

According to the present theory, this possible way of talking about the ground by reference to properties of the sky is very much like the way we actually talk about ways of looking by reference to properties of objects that have them. There are properties of ways of looking that we cannot attribute directly, because we have no expressions that stand for them. If x looks red, for example, then the way that x looks has a certain property, but we have no expression that stands for this property. If we did, say ‘is P’, then we could say that x looks red by saying that the way that x looks is P. But we don’t. Rather, we say that x looks red. Even though this form of words might suggest it, we are *not* saying that x is represented as being red: we are saying that the way that x looks is P, and even though this may be a sign that x is red, it does not *represent* that x is red. I suggest that if we had a stock of expressions sufficiently rich that we could attribute properties to ways of looking by using sentences of the form ‘w is P’ then we would not mistakenly think that the visual experience of an object represents it to be a certain way.

This also provides an answer to another pressing question: if S’s experience represents the way that x looks rather than the way that it is, then how does it make sense to ask whether or not her experience is veridical? On the traditional view the answer is clear: if x looks red to S then the content of her experience is that x is red, and

the experience is veridical just in case *x* is red. This allows for both veridical and non-veridical experience, because *x* being red and *x* not being red are both compatible with *x* looking red to *S*. On the view here it is not so clear: if *x* looks red to *S* then the content of her experience is the way that *x* looks to *S*, and we can't say that the experience is veridical just in case that *is* the way that *x* looks to *S*, for then we could never truly say that *x* looks red to *S* but her experience is non-veridical: that not being the way that *x* looks to *S* is not compatible with that being the way that *x* looks to *S*.¹⁵ But the answer, again, lies in the peculiar way in which we talk about ways of looking. Because we attribute properties to the way that *x* looks by reference to properties that it may or may not have, then it is natural to ask whether or not *x* does have those properties. To say that *x* looks red to *S* is to say that the way *x* looks has a certain property, one that is somehow connected but not identical to the property of being red, and it is in virtue of this connection that there is a sensible notion of veridicality. But note: the veridicality of an experience is not, on this view, the truth of its representational content.

According to the theory developed here, then, Peacocke's scenarios are not a suitable basis for any theory of the content of visual experience, conceptual or nonconceptual, because visual experience does not have content of that kind: the content of a visual experience is, rather, a way of looking. But perhaps ways of looking can play the role that Peacocke wants scenarios to play. Like scenarios, they are not 'built up' from concepts, and there is no requirement that their specification be limited to concepts possessed by the subject of the experience. Furthermore, I suggest that they are equally well-placed to explain three key features of the content of visual experience that Peacocke appeals to: its fine-grained structure, its analogue and unit-free nature, and its overlap with the content of experiences in other sense modalities. I won't argue for these claims here.

¹⁵ If there is no real notion of the truth or falsity of an experience's content, then perhaps it does not make sense to talk of the experience having content. Perhaps rather than claiming that the content of visual experience is not that the world is thus and so, I should make the stronger claim that it has no content at all.

Chapter 3

Knowing About How Things Look

In this third chapter I will defend the theory developed in Chapters 1 and 2 against one possible objection: that it severs an intuitively plausible connection between how things look and what we can know about how things look. In particular, it provides no guarantee that we can know how things look just by looking at them, a guarantee that we often take ourselves to have. I will give independent reasons for thinking that we have no such guarantee, and I will show that by allowing that sometimes we cannot come to know how things look just by looking at them we can view so-called *phenomenal* and *non-phenomenal* sorites paradoxes as instances of a single, common phenomenon.

3.1 Some Plausible Principles

It is an intuitively attractive thought that to say that x *looks the same as* y in respect of R to observer S at time t is to say that S lacks a certain discriminatory ability.

Following Williamson (1990), I will take it that to say that S lacks a certain discriminatory ability is to say that S cannot come to know, in the relevant kind of way, that two things are distinct. Which things? I am happy to say: *the way that x looks in respect of R to S at t* and *the way that y looks in respect of R to S at t* , because I believe that there are such things. Then to say that x looks the same as y in respect of R to S at t is to say that S cannot come to know, in the relevant kind of way, that the way that x looks in respect of R to her at t is distinct from the way that y looks in respect of R to her at t . But not everyone will accept that there are such things. So instead I will take it that the discrimination in question is between x and y themselves, and that to say that x looks the same as y in respect of R to S at t is to say that S cannot come to know, in the relevant kind of way, that x and y are distinct.

What is the relevant kind of way? S comes to know that x and y are distinct by coming to know that one has a property that the other does not. But not all properties are relevant in this case. Coming to know that x *is here* but that y is not, and thereby coming to know that x is distinct from y, does not count as discriminating x from y in the relevant kind of way, because it may still be that x looks the same as y. Nor does coming to know that x *is square* but that y is not, nor that x *weighs 10kg* but that y does not, nor that x *likes chocolate* but that y does not, and so on. For each of these properties it is possible that only one of x and y has it, and yet that x and y still look the same. The relevant properties in this case are not ones that objects do or do not have in virtue of the way they *are*, but ones that they do or do not have in virtue of the way they *look*, properties that we attribute by using *phenomenal* looks sentences and express by using predicates of the form ‘looks F’ where ‘F’ refers to some element of some respect of ways of looking (predicates like ‘looks red’, ‘looks square’, and so on). So the relevant kind of way for S to come to know that x and y are distinct is for S to come to know that for some R and some F, x looks F in respect of R to S at t but y does not.¹⁶

But that is not quite enough. If there is no way that x and y look to S at t then there is no F such that x looks F to S at t, so there is no F such that x looks F to S at t but y does not, so S cannot come to know that there is. But it is not the case that x looks the same as y to S. So we should add that x and y must look some way to S at t. But that is not quite enough either. Suppose that there is some way that x and y look to S, that S *cannot* come to know, just by looking at x and y, that there is an F such that x looks F in respect of R to her but y does not, but that she *can* come to know that there is when she looks at some third thing z as well. The intuitively attractive thought under consideration here is that in this case x still *does* look the same as y to S. If she cannot come to know, *just by looking at x and y*, that there is an F such that x looks F to her but y does not, then x and y look the same to her.

When S lacks the discriminatory ability in question here it is variously said that x *matches* y (in respect of R to S at t), or that x *is phenomenally indiscriminable from* y, or that x *is phenomenally indistinguishable from* y. And if S has that discriminatory ability it is variously said that x *does not match* y (in respect of R to S at t), or that x *is*

¹⁶ Unless otherwise stated, in this chapter I shall only be concerned with phenomenal ‘looks’ sentences - sentences that express propositions of the form F1 using sentences of the form ‘x looks F’ where ‘F’ refers to some element of some respect of ways of looking.

phenomenally discriminable from y, or that *x is phenomenally distinguishable from y*.
So I will take it that:

(Matching) To say that *x matches y* in respect of R to S at t (or *is phenomenally indiscriminable from*, or *is phenomenally indistinguishable from*) is to say that x and y both look some (possibly different) way in respect of R to observer S at time t, and that S *cannot* come to know, just by looking at x and y at t, that there is an F such that x looks F in respect of R to S at t but y does not. To say that *x does not match y* in respect of R to S at t (or *is phenomenally discriminable from*, or *is phenomenally distinguishable from*) is to say that x and y both look some way in respect of R to observer S at time t, and that S *can* come to know, just by looking at x and y at t, that there is such an F.

The qualification ‘phenomenally’ in ‘is phenomenally indiscriminable from’ (and in ‘is phenomenally discriminable from’) is important. ‘Matches’ and the unqualified ‘is indiscriminable from’ stand for different relations: If x and y are colour patches that I cannot see, then I *would* assert that x is indiscriminable from y in respect of colour (by me, now), but would *not* assert that x matches y in respect of colour. Indiscriminability, simpliciter, is a more widespread phenomenon than *matching*, and only a particular kind of indiscriminability is relevant here. The word ‘phenomenal’ is intended to specify which kind that is. Similar remarks hold for the qualification ‘phenomenally’ in ‘is phenomenally indistinguishable from’ (and in ‘is phenomenally distinguishable from’).

As I am taking it, ‘matches’ stands for a relation that obtains between the *objects* that look a certain way (in a particular respect, to a particular observer, at a particular time), rather than between the *ways* that those objects look, or between the *elements of the respect* in which those objects look (colours, shapes, and so on). Graff does the same. Goodman (1951) and Peacocke (1981) use it to stand for a relation that obtains between phenomenal qualities of presentations of the objects. Williamson (1990) and

Clark (1993) use it to stand for a relation that obtains between the presentations themselves, in virtue of their qualities. The difference won't matter here.

With the relation of *matching* thus understood, I will take it that the intuitively attractive thought we are considering can be put this way:

(Principle1) To say that *x looks the same as y* in respect of R to S at t is to say that *x matches y* in respect of R to S at t.

Here is another intuitively attractive thought:

(Principle2) *x looks the same as y* in respect of R to S at t if and only if for all F, if *x looks F* in respect of R to S at t then *y looks F* in respect of R to S at t.

The thought is not just that Principle 2 is true, but that it is a conceptual truth – it is true in virtue of what ‘looks F’ means: anyone who claims, for example, (a) that *x looks the same as y* in respect of colour to her but that, for some colour F, *x looks F* to her whereas *y* does not, or (b) that there is no colour F such that *x looks F* to her whereas *y* does not, but that *x looks different from y*, would seem not just to be mistaken, but to misunderstand what it is to look F.

This is a thought about the meaning of *observational* predicates that is similar to one about the meaning of *vague* predicates. It is attractive to think that it is part of the meaning of a vague predicate such as ‘is tall’ that if *x* and *y* differ in height by less than 1 millimetre, then if *x* is tall then *y* is tall: anyone who claimed otherwise would seem not just to be mistaken but to misunderstand the meaning of ‘is tall’, and more generally about what it is for a predicate to be vague. So too, it is attractive to think that it is part of the meaning of an *observational* predicate such as ‘looks F’ that if *x looks the same as y* in respect of R to me now, then if *x looks F* then *y looks F*: anyone who claimed otherwise would seem not just to be mistaken but to misunderstand the meaning of ‘looks red’, and more generally about what it is for a predicate to be observational.

Putting together Matching, Principle1 and Principle2 we get the following:

(Observationality) If x and y both look some (possibly different) way in respect of R to observer S at time t, then there is an F such that x looks F in respect of R to S at t but y does not, if and only if S can come to know that there is, just by looking at x and y at t.

Observationality can be thought as capturing a third intuitively attractive thought – that we can know how things look just by looking at them. Indeed, it is attractive to think that not only is Observationality true, but that it too is a conceptual truth – it is true in virtue of the meaning of ‘looks F’: if S were to claim that she cannot come to know, even by looking at x and y, that there is an F such that x looks F to her but y does not, and yet that it is still may be that there is such an F, she would seem not just to be mistaken, but to misunderstand what it is to look a certain way to her. After all, if there is anything different about x and y that she cannot come to know of just by *looking* at them, then surely it cannot be anything about the way that they *look* to her.

Observationality, when it is taken as a conceptual claim, is at odds with the theory developed in Chapters 1 and 2. According to that theory, to say that x looks F in respect of R to S at t is *not* to say anything about what S can or cannot come to know about x or its identity or non-identity with anything else; it is to say something about a certain thing: the way that x looks to S at t. The theory provides no *conceptual* guarantee that if there is an F such that x looks F in respect of R to S at t but y does not then S can come to know that there is, or, contrapositively, that if S cannot come to know that there is such an F then there is no such F. If x and y both look a certain way in respect of R to S at t then they do so conceptually independently of what S can come to know, and it may be that x, but not y, looks F to S at t, even though S cannot come to know that it does. If that theory is correct then it follows that Observationality is not a conceptual truth. But if Observationality is not a conceptual truth, then it follows that either Principle1 is false or Principle2 is not a conceptual truth, since together these imply that Observationality is a conceptual truth.

To avoid this being seen as a *reductio ad absurdum* of that theory, my aim in this chapter is to give independent reasons for thinking that both Observationality and Principle1 are false. As I showed in Section 2.6.2, Principle2 is a consequence of the

theory, so I do not want to deny its truth. But Principle1 is not. In fact, it too is at odds with the theory. According to the theory, to say that x looks the same as y in respect of R to S at t is *not* to say anything about what S can or cannot come to know about the identity of x and y. Rather, it is to say that two things are identical: the way that x looks in respect of R to S at t, and the way that y looks in respect of R to S at t. So if that theory is correct then Principle1 is false.

In Section 3.2 I will give a variety of arguments for the falsity of Principle1. In Section 3.3 I will offer reasons why Principle1 and Observationality might seem intuitively plausible, even if the first is not a truth and the second is not a conceptual truth. In Section 3.4 I will discuss an apparent problem for Observationality, by considering sorites-style paradoxes that are thought to arise if Observationality is true. Finally, in Section 3.5 I will argue that by denying the truth of Observationality it is possible to resolve these sorites-style paradoxes in a way that is akin to those that arise from vagueness.

3.2 Problems for Principle1

In this section I will raise some problems for Principle1.

3.2.1 Mismatch with Other Ways

To say that x *walks the same as* y, that x *dresses the same as* y, that x *is the same as* y, that x *talks the same as* y, that x *reacts the same as* y, that x *flies the same as* y, and so on, is not to say, we take it, that some observer cannot come to know in the relevant kind of way that x is distinct from y. So to claim that to say that x *looks the same as* y is to say that some observer cannot come to know in the relevant kind of way that x is distinct from y is to claim that there is some part of the meaning of ‘looks the same as’ that is not part of the meaning of ‘walks the same as’, ‘dresses the same as’, ‘talks the same as’, and so on.

But if asked what is the difference between the meaning of ‘walks the same as’ and ‘talks the same as’, it is natural to say that one is about the way things walk, whereas the other is about the way things talk: the first says something about the way they walk,

whereas the second says it (i.e. that same thing) about the way they talk. So too, if asked what is the difference between the meaning of ‘walks the same as’ and ‘looks the same as’, it is natural to say that one is about the way things walk, whereas the other is about the way things look: the first says something about the way they walk, whereas the second says it (i.e. that same thing) about the way they look. It is not natural to mention any difference between the meaning of ‘walks the same as’ and ‘looks the same as’ that is not mentioned about the meaning of ‘walks the same as’ and ‘talks the same as’. But it would be, if there were such a difference.

Compare this with what we might say about ‘starts before’, ‘finishes before’ and ‘stands before’. There is something in common about the meaning of all three – they each say that things stand in a certain spatio-temporal relation. But there is part of the meaning of ‘stands before’ that is not part of the meaning of the other two – for A to stand before B not only must she be in a certain spatio-temporal location relative to B, but she must also be facing B. And we are inclined to mention this difference when asked. If asked what is the difference between the meaning of ‘starts before’ and ‘finishes before’, it is natural to say that one is about when things start whereas the other is about when they finish: the first says something about the times that they start, whereas the second says it (i.e. that same thing) about the times that they finish. If asked what is the difference between the meaning of ‘starts before’ and ‘stands before’, it is natural to say that one is about when two people start whereas the other is about where they stand *and* how they are facing: the first says something about the times that they start, whereas the second says it (i.e. that same thing) about where they stand, *and also* something about where they are facing. We mention this extra difference between ‘starts before’ and ‘stands before’ because there is one, so if we do not mention any extra difference between ‘walks the same’ and ‘looks the same’ then the natural explanation is that there isn’t one.

If Principle1 is true then there is a certain difference between the meaning of ‘looks the same as’ and the meaning of ‘walks the same as’, ‘dresses the same as’, and so on. But much of the way we talk is as though there is no such difference.

3.2.2 Exactly?

Bill and Ben are discussing colour patches x and y that only Ben can see. Here is part of their conversation:

Bill: How do x and y look in respect of colour?

Ben: They look the same.

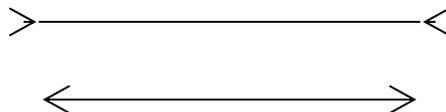
Bill: You mean they look *exactly* the same?

Ben: Well, I don't know if they look *exactly* the same. But they're close enough that I can't tell any difference.

With his last comments Ben is saying, I take it, at least these two things: (a) that it might be false that x looks the same as y, and (b) that, even so, he cannot perceptually discriminate x from y. Thus Ben is allowing that in this case 'x looks the same as y' expresses a proposition that is false whereas 'x matches y' expresses a proposition that is true. This does not show that there actually *can* be cases in which they express propositions with different truth values. But it does show, I think, that Principle1 is false. For if to say that x looks the same as y is to say that x matches y, as Principle1 claims, then if Ben understands these expressions (which we can assume, or add to the story) then he ought to think that they express the same proposition, and so to not allow that they can express propositions with different truth values. But he does. So to say that x looks the same as y cannot be to say that x matches y.

3.2.3 The Muller-Lyer Illusion

It is widely believed that in the Muller-Lyer illusion, the top line does not look the same as the bottom line in respect of length (it looks longer):



If that is right, and if Principle1 is right, then that is to believe that an observer can come to know, just by looking at the two lines, that there is an F such that the top line looks F in respect of length but the bottom line does not. And we often think that we can – that we can *see* that there is such an F, and thereby come to know it, just by looking at the lines.

But I, for one, only take myself to be seeing that there is such an F when I look at the diagram in a particular way – a way in which I am not looking *closely* at the diagram, and, in particular, not looking closely at either line. If I *do* look closely at the lines – for example, by looking at their left-hand ends and imagining a vertical line joining those ends, and by looking at their right-hand ends and imagining a vertical line joining those ends – then I am no longer inclined to think that I can see that there is such an F. Looking too closely at the lines undermines my belief that I can see that there is such an F.

But it is familiar that how something looks to an observer depends upon how the observer is looking at it: a plate that looks circular when viewed square-on may look oval when viewed from an angle; a book that looks red when located in the centre of one's visual field may look blue when located in the periphery of that field, and so on. So one natural response to my comments about the Muller-Lyer illusion is this: that even if I cannot see, when I look at the lines *closely*, that there is an F such that, when I look at the lines not-closely, the top line looks F in respect of length but the bottom line does not, it may still be the case that I can see, when I look at the lines *not-closely*, that there is such an F when I look at them not-closely, because it may well be the case that the way the lines look to me when I look at them closely is different from the way the lines look to me when I look at them not-closely.

I'll grant that that is right. But the point is this: if I take it that I can see that there is an F such that the top line looks F in respect of length but the bottom line does not, *only* if I look at the lines not-closely, and that I cannot see that there is such an F when I look at them closely, then can I *really* take myself to be able to come to know that there is such an F? If I cannot look closely at the lines to check that there is such an F when I look not-closely, then can I really *know* that there is? I take it that by looking not-closely I can see that there is such an F, but if I can't look more closely to *make sure* that I can see it, then can I *be* sure that I really can see it?

Perhaps these considerations show that we *can't* come to know, just by looking at the two lines, that there is an F such that the top line looks F in respect of length but the bottom line does not (when looked at in the right way). If that is right, and if Principle1 is right, then despite what is widely believed, the top line *does* look the same as the bottom line in respect of length. There would be a tension between Principle1 and the

claim that in the Muller-Lyer diagram the top line looks longer than the bottom line, and the extent to which we uphold the latter is the extent to which we must reject Principle1. If we want to uphold Principle1 then we must deny that the top line looks longer.

Perhaps that's not such a hard pill to swallow – the Muller-Lyer illusion is an exceptional case, and if it turns out that in order to uphold Principle1 we have to accept that the lines do after all look the same in respect of length, then so be it. But if the argument does succeed then there are many harder pills to swallow. Suppose that *x* and *y* are colour patches side-by-side on a wall, and that I take myself to be able to see, quite clearly, that *x* does not look the same as *y* in respect of colour. If Principle1 is right, then that is to take myself to be able to see that there is an *F* such that *x* looks *F* in respect of colour but *y* does not. But I think it can be argued, again, that there is a limit to how closely I can look at *x* while still seeing that there is such an *F*. And the argument is quite simple: to see that there is such an *F*, I must be able to see both *x* and *y*, and so I must be looking at both *x* and *y*; but having to keep looking at *y* imposes a limit on how closely I can look at *x*; I cannot, for instance, give *x* my *full* attention; but if there is a limit to how closely I can look at *x*, then there is a limit to how well I can make sure that there actually *is* such an *F*, which casts doubt upon whether I can really *know* that there is such an *F*. If these considerations show that I *can't* come to know, just by looking at *x* and *y*, that there is an *F* such that *x* looks *F* in respect of colour but *y* does not, and if Principle1 is right, then despite what I firmly believe, *x* *does* look the same as *y* in respect of colour. And if these considerations show that, then I think that similar considerations will show, for *all* objects *x* and *y* and respects *R*, that I can't come to know, just by looking at *x* and *y*, that there is an *F* such that *x* looks *F* in respect of *R* but *y* does not. If that is right, and if Principle1 is right, then *no* two objects look different to me. Anyone who wants to uphold Principle1 must deny that *any* two things look different.

If the arguments that I have been putting forward here really do conclude that for no *x*, *y*, *S* and *R* can *S* come to know that there is an *F* such that *x* looks *F* in respect of *R* to her but *y* does not, then perhaps this shows that they must be flawed: such a skeptical result is unacceptable. I'll grant that – I'll grant that we do indeed have such knowledge, and that the arguments do not succeed in showing otherwise. But I think they can still be used to show that Principle1 is false. Consider the Muller-Lyer illusion

again. Suppose I argue in the way above to make you believe, rightly or wrongly, that you *cannot* come to know, just by looking at the two lines, that there is such an F such that the top line looks F in respect of length but the bottom line does not. You may well nevertheless still *believe* that there is, and so that the top line looks longer, even though you believe that you can't come to know that there is. But if Principle1 is right, then coming to believe that you cannot come to know that there is such an F just *is* coming to believe that the lines look the same in respect of length. So if Principle1 is right, then you ought not to believe *both* that you cannot come to know that there is such an F *and* that the top line does not look the same as the bottom line in respect of length. But that is something you may well do, and not because you misunderstand 'looks the same as' (we can assume, or add to the story, that you do understand it). Of course, it does not follow from this that you are ever right to hold such beliefs. But because we can hold such beliefs without being confused in any way, believing that we cannot come to know that there is such an F cannot be believing that the lines look the same in respect of length. Principle1 cannot be true.

3.3 The Plausibility of Principle1 and Observationality

If I believe that colour patches x and y match in respect of colour to me now, then I might well say that x looks the same as y (in respect of colour to me now). Anyone who upholds Principle1 can easily explain why: to say that x looks the same as y just is to say that x matches y, so all that I am doing is reporting my belief. But according to the theory developed in Chapters 1 and 2, to say that x looks the same as y is *not* to say this at all. Rather, it is to say that the way that x looks (in respect of colour from some aspect a now) is identical to the way that y looks (in respect of colour from some aspect b now). But, as I argued in Section 3.2.2, I might well say that x looks the same as y, even though I do *not* believe that the way that x looks is identical to the way that y looks. Moreover, I might well say that x looks the same as y even though I believe that the way that x looks is *not* identical to the way that y looks. That is, I might well say something that I believe is false. If I want to reject Principle1 then I need another way to explain this behaviour.

It is commonplace, on the face of it at least, for us to say things that we believe are false. You ask me what time it is. I look at my watch and see that it is 1:57 p.m., but I tell you that it is 2:00 p.m. You ask me the population of Australia. I know that it is 18.1 million, but I tell you that it is 18 million. You ask me how Bill and Ben compare in respect of height. I can see that Bill is slightly taller than Ben, but I tell you that they are the same height. Well, at least there are some contexts in which I might tell you these things. But there are also contexts in which I would not. If you had asked me the time because you were setting your watch, then I would have told you that it is 1:57 p.m. rather than 2:00 p.m. If you had asked me the population of Australia because you are compiling an encyclopaedia, then I would have told you that it is 18.1 million rather than 18 million. If you had asked me how Bill and Ben compare in respect of height because you want to find the tallest person in the room, then I would have told you that Bill is taller than Ben rather than the same height. So context must play a role in any explanation of why I tell you what I do.

One explanation is this. The meaning of ‘it is 2:00 p.m.’ varies from context to context, depending upon the standards of precision set by those contexts. When I utter ‘it is 2:00 p.m.’ in the first context I say something true, and I believe that I do, even though it is actually 1:57 p.m., and even though I believe that it is. But if I were to utter ‘it is 2:00 p.m.’ in the second context I would say something false, and I believe that I would, because it is actually 1:57 p.m., and I believe that it is. So despite appearances, I am *not* saying things that I believe are false. I am saying things that I believe are true. According to this explanation, the context helps determine what my utterance expresses.¹⁷

Another explanation is this. The meaning of ‘it is 2:00 p.m.’ does not vary from context to context. When I utter ‘it is 2:00 p.m.’ in the first context I say something false, and I believe that I do, because it is actually 1:57 p.m., and I believe that it is. But the context is such that *it will do you no harm* to take it to be 2:00 p.m. rather than 1:57 p.m., and it is more expedient for me to do so, so I tell you that instead. Suppose you are asking me the time because you have to go somewhere and you want to know whether or not it is time to leave. I know that telling you it is 1:57 p.m. and telling you it is 2:00 p.m. will both lead you to leave, so I tell you that it is 2:00 p.m. because it will do you

¹⁷ See Lewis (1979) for an account along these lines, especially Example 5: Vagueness.

no harm, and because it is more expedient to do so. But suppose you are asking me because you want to set your watch to the correct time. Then I know that telling you it is 1:57 p.m. and telling you it is 2:00 p.m. will lead you to take different courses of action, the first in which you set your watch to the correct time, the second in which you set your watch to an incorrect time. Telling you in this context that it is 2:00 p.m., or any time other than 1:57 p.m., will do you harm. So I don't. According to this explanation, the context helps determine whether or not my utterance does you any harm.

I want to use an explanation of this second kind to suggest why I might say that *x* looks the same as *y* (i.e. that the way that *x* looks is identical to the way that *y* looks), even though I do not believe that it does. If *x* and *y* match in respect of colour for me, then even if the way that *x* looks in respect of colour and the way that *y* looks in respect of colour are not identical, they are nevertheless similar enough that I cannot come to know that they are distinct. If I say that *x* looks the same as *y* when all I believe is that they match, then I do so because I believe that the way that *x* looks and the way that *y* looks are similar enough that taking them to be identical is expedient and will do no harm.

I think we can use this to explain why anyone might think that there is a closer connection between the way something *looks* and what we can know about it than there is between the way something *is* (or *walks*, or *talks*, and so on) and what we can know about it. It is easy to imagine a context in which I cannot discriminate between *x* and *y* in respect of how they *are*, but in which it might do harm to take it that *x* is the same as *y*. Suppose, for example, that I am baking a cake and need to use self-raising rather than plain flour. The flour in one container looks and tastes the same as the flour in the other, and neither container is labelled. So I cannot discriminate one flour from the other in respect of whether or not it is plain or self-raising. Nevertheless, it might do me harm to take them to *be* the same, because if I use plain flour then my cake will not rise. But it is not so easy to imagine a context in which I cannot discriminate between *x* and *y* in respect of how they *look*, but in which it might do harm to take it that *x* looks the same as *y*. If I cannot come to know, just by looking at *x* and *y*, that there is an *F* such that *x* looks *F* to me but *y* does not, then what harm can it do to take it that there is no such *F*? After all, any difference between the way that *x* looks and the way that *y* looks can only

be causally efficacious, for instance, if I can tell that there is a difference, so if I can't then I might as well take it that there isn't.¹⁸

And this is why, I suggest, it is attractive to think that 'looks the same as' just means 'matches'. Whenever x and y match it will do no harm, we believe, to take it that they look the same, even though they may not. It is appropriate, and expedient, to say that x looks the same as y, so we say that x looks the same as y. Because we utter 'x looks the same as y' whenever x and y match, we are mistakenly led to think that it *means* that x and y match. But it does not.

3.4 The Problem of the Non-Transitivity of *Matching*

Matching is a kind of indiscriminability, and indiscriminability is a non-transitive relation. I can see that the person on the left is male and that the person on the right is female, but I cannot see the person in the middle. Thus in the circumstances I *cannot* discriminate the person on the left from the person in the middle in respect of sex, nor the person in the middle from the person on the right in respect of sex, but I *can* discriminate the person on the left from the person on the right in respect of sex. But it may still be that particular *kinds* of indiscriminability are transitive. After all, *being due north or due east of* is a non-transitive relation, but *being due north of*, one of its particular kinds, is transitive. So just because indiscriminability is non-transitive, it does not follow that *matching*, one of its particular kinds, is non-transitive.

Nevertheless, it is widely believed that it is. Goodman (1951, p. 196) says that "[a]lthough two qualia q and r exactly match, there may be a third quale s that matches one but not the other". Armstrong (1968, p. 218) asks us to consider three samples of cloth, A, B and C, such that "A and B are perceptually completely indistinguishable in respect of colour, and B and C are perceptually completely indistinguishable in respect of colour... [but] A and C can be perceptually distinguished from each other in this respect". Dummett (1975, pp. 316-7) says that "... there does not seem to be any doubt that there is such a relation as non-discriminable difference (of position, colour, etc.), and that it is non-transitive". Wright (1975), Peacocke (1981), and Everett (1996) all assume the non-transitivity of *matching* in their arguments.

¹⁸ I will argue in Section 3.5 that there is, surprisingly, at least one context in which it *does* do harm.

So the claim is this:

(NonTrans) *Matching* (*phenomenal indiscriminability, phenomenal indistinguishability*) is a non-transitive relation.

That is, it is possible for there to be x , y and z that all look some way in respect of R to S at t , such that S cannot come to know, just by looking at x and y at t , that there is an F such that x looks F (in respect of R to S at t) but that y does not, and S cannot come to know, just by looking at y and z at t , that there is an F such that y looks F but that z does not, but S *can* come to know, just by looking at x and z at t , that there is an F such that x looks F but that z does not.

It is, indeed, easy to convince oneself that *matching* is non-transitive, by imagining examples of what are sometimes called *phenomenal sorites series*: perhaps a long series of colour patches p_1, \dots, p_n such that the left-most patch p_1 looks red (to me, now), such that each patch p_k matches the patch on its right p_{k+1} in respect of colour, but such that the right-most patch p_n looks blue and does not match the left-most patch p_1 in respect of colour; or perhaps a long series of shapes s_1, \dots, s_n such that the left-most shape s_1 looks circular (to me, now), such that each shape s_k matches the shape on its right s_{k+1} in respect of shape, but such that the right-most shape s_n looks square and does not match the left-most shape s_1 in respect of shape.

If *matching* is indeed non-transitive, then there is a problem for Observationality. For suppose that x matches y (in respect of R to S at t), y matches z , but x does not match z . Since x does not match z , according to Observationality there is an F such that x looks F (in respect of R to S at t) but z does not.¹⁹ Since x matches y , according to Observationality there is no F such that x looks F but y does not. So since x looks F , it must be that y looks F also. Since y matches z , according to Observationality there is no F such that y looks F but z does not. So since y looks F , it must be that z looks F also. But z does not look F . If *matching* is non-transitive then Observationality leads to contradiction.

This problem is not due to any vagueness that there might be in 'F'. If it were, then we should be able to avoid it by precisifying 'F'. Without loss of generality, take 'F' to

¹⁹ In fact, this follows from the principle of Matching and the factivity of knowledge.

name the colour red. Let's precisify it by defining it to name that range of colours between the maximally specific colours r and r_+ inclusive. If x looks any way in respect of colour (to S at t), then there is a maximally specific colour C such that x looks C , even if S cannot say which colour that is. Since x looks red, C is some colour in the range from r to r_+ (inclusive). According to Observationality, since x matches y in respect of colour and since x looks C in respect of colour, y also looks C in respect of colour, and therefore y looks red. And according to Observationality, since y matches z in respect of colour and since y looks C in respect of colour, z also looks C in respect of colour, and therefore z looks red. So precisifying 'red' does not remove the problem, so its source cannot be any vagueness that there might be in 'red'. More generally, no matter what 'F' names precisifying 'F' will not remove the problem, so the source of the problem cannot be any vagueness that there might be in 'F'.

We have been discussing circumstances in which x and y both look some way in respect of R to S at t , and in which there either is or is not an F such that x looks F in respect of R to S at t but y does not. If there is not, let's say that x *has the same look as* y (in respect of R to S at t), and if there is, let's say that x *has a different look from* y . Then Observationality is the claim that x has the same look as y if and only if x matches y .²⁰ That is, that a necessary and sufficient condition for the *has the same look as* relation to hold between x and y is that the *matching* relation holds between x and y . But *has the same look as* is a transitive relation. For suppose that it were not. Then there would be x , y and z such that x has the same look as y , y has the same look as z , but x has a different look from z . Since x has a different look from z , there is an F such that x looks F (in respect of R to S at t) but z does not. Since x looks F it must be that y looks F also, otherwise it would not be true that x has the same look as y . Since y looks F it must also be that z looks F also, otherwise it would not be true that y has the same look as z . But z does not look F . So there are no such x , y and z , and *has the same look as* is transitive.

Since *has the same look as* is transitive, if *matching* is not transitive then it cannot be a necessary and sufficient condition for the former, and so Observationality must be false. In that case we might try modifying the formulation of Observationality by replacing the non-transitive relation of *matching* by some *transitive* relation. This is

²⁰ Actually, that is what its contrapositive claims, but the two are necessarily equivalent.

essentially what Goodman did, and in the spirit (but not the letter) of his approach we might modify the principle to this (where the subscript ‘G’ is for ‘Goodman’):

(Observationality_G) If x and y both look some way in respect of R to observer S at time t , then x has the same look as y in respect of R to S at t if and only if for all z , if x matches z in respect of R for S at t then y matches z in respect of R for S at t .

The relation on the right-hand side of the biconditional here is transitive, even if *matching* is not. For call it *matching*_G, so that x matches_G y if and only if for all z , if x matches z (in respect of R for S at t) then y matches z . Now suppose that *matching*_G is not transitive, so that there are x_1 , x_2 and x_3 such that x_1 matches_G x_2 , x_2 matches_G x_3 , but x_1 does not match_G x_3 . Since x_1 does not match_G x_3 , there is some z such that x_1 matches z but x_3 does not match z . But x_1 matches_G x_2 , so if x_1 matches z then x_2 matches z . But x_2 matches_G x_3 , so if x_2 matches z then x_3 matches z . But x_3 does not match z . So we have a contradiction, and *matching*_G must be transitive.

Unlike Observationality, Observationality_G is not threatened by the supposed non-transitivity of *matching*. But, unlike Observationality, it does not capture the intuitively attractive thought that we would like it to: that we can know how things look *just by looking at them*. If x and y both look some way in respect of R to me now, and I cannot come to know, just by looking at x and y , that there is an F such that x looks F to me now but y does not, then, according to Observationality it follows that there is no such F , regardless of whether or not there are other things that look any way in respect of R to me now, and, if there are, regardless of whether or not they match x and/or y . But according to Observationality_G, even if I cannot come to know, just by looking at x and y , that there is such an F , it may still be that there is – there may be something about the way that x and y look that I cannot come to know of just by looking at x and y . In fact, there *is*, if x and y are two of three objects for which the transitivity of *matching* fails, for then there is a third object z that looks some way to me in respect of R , such that z matches only one of x and y in respect of R , from which it follows, according to Observationality_G, that the way that x looks to me is different from the way that y looks to me. So the non-transitivity of *matching* poses a problem for Observationality_G as

well: that it is possible for there to be objects x and y such that for some F , x looks F to me but y does not, even though I cannot come to know that, just by looking at x and y . The problem is that not only is this *not* intuitively attractive, but it is intuitively *unattractive*. Graff (2001) puts it this way:

The opponent of transitivity is thus committed to the possibility that A and B might be perceptually indistinguishable yet have different looks. But what I have difficulty understanding is, if it is really the look of the two which is different, why cannot this come across, so to speak, by looking just at *them*? Why must one look at some third thing to see that the looks of A and B are distinct? (pp. 915-6).

... I find Goodman's distinction between looking the same and having the same look to be barely coherent... (p. 916)

So by retreating from Observationality to Observationality_G we can avoid the *logical* problem posed by the purported non-transitivity of *matching*, but in so doing we face a *conceptual* problem: that Observationality_G does not express the intuitively attractive thought that we'd like it to.

Another response to the problem posed by phenomenal sorites series is to argue, as do Raffman (2000) and Graff (2001), that *matching* is actually *transitive*, and hence that such series are impossible. If they are impossible, then we can avoid the paradox that they are thought to generate while holding onto the truth of Observationality.

But the theory developed in Chapters 1 and 2 opens up the possibility of another response. As I mentioned in Section 3.1, it is attractive to think that Observationality is a conceptual truth, and therefore a necessary truth. If it is, then in response to the problem posed by phenomenal sorites series we cannot, without confusion, deny the truth of Observationality. But I have paved the way to doing just that by arguing against the truth of Principle1 (Section 3.2), and by suggesting why Principle1 and Observationality might be so intuitively attractive despite being false (Section 3.3). If we deny the truth of Observationality then we have a solution to the problem posed in this section which allows us to uphold the non-transitivity of *matching*. Moreover, it opens up the possibility of giving a uniform response to both phenomenal and non-

phenomenal sorites paradoxes. My aim in the next and final section is to show how it does this.

3.5 Phenomenal and Non-Phenomenal Sorites Series

In ‘Phenomenal Continua and the Sorites’, Delia Graff draws a distinction between *non-phenomenal* and *phenomenal* sorites arguments (Section 1, pp. 907-8). As an example of the former she gives:

Argument A

- Premise 1: Any seven-foot tall man is tall.
Premise 2: Any man one-hundredth of an inch shorter than a tall man (the inductive premise) is tall.
Conclusion: Therefore, any three-foot man is tall.

As an example of the latter she asks us to “consider a ‘sorites series’ of thirty colour patches, uniformly changing in appearance from one that looks red on the left toward one that looks yellow on the right, while each patch looks the same in respect of colour as its successor in the series” (p. 907), about which it can be argued:

Argument B

- Premise 1: The first patch looks red.
Premise 2: If any two patches look the same in respect of colour, (the inductive premise) then if one looks red so does the other.
Conclusion: The thirtieth patch looks red.

What makes this a *phenomenal* rather than *non-phenomenal* sorites argument, she explains, is that it contains *observational* relational terms and predicates like ‘looks the same’ and ‘looks red’, rather than *non-observational* ones like ‘is shorter than’ and ‘is tall’, where a predicate is observational just in case its applicability to an object depends only on the way that object *appears* (looks, sounds, tastes, etc.), not upon the way that object *is* (and where, presumably, a relational term is observational just in case its

applicability to a pair of objects depends only on the way those objects appear, perhaps separately, perhaps together).

These arguments are paradoxical because apparently valid reasoning takes us from apparently true premises to an apparently false conclusion. Graff's opinion is that in both cases the reasoning is indeed valid and the conclusions are indeed false, but otherwise she locates the blame differently. In the first case she thinks that the inductive premise is to blame, it being "false that if any two men differ in height by one-hundredth of an inch, then if one is tall so is the other". But in the second she thinks that it is not to blame, it being "*true* that if any two colour patches look the same, then if one looks red so does the other". She places the blame in this case on the series of colour patches that we are being asked to consider, claiming that the existence of such a series is a logical impossibility. She seems to think that these are different resolutions, because she asks, "Why the difference in attitude?" (p. 908). And she seems to think that what requires them to be resolved differently is that one argument is non-phenomenal and the other phenomenal. She says:

If two men differ in height by even one-hundredth of an inch, then they differ in a respect that is relevant for the applicability of 'tall'. But if two colour patches look the same (not just similar, but the same) in respect of colour, then they do not differ, on the face of it at least, in any respect relevant for the applicability of 'looks red'. (p. 908)

According to the theory developed in Chapters 1 and 2, Graff is right to claim that the existence of such a series of colour patches is a logical impossibility. It follows from that theory that *looks the same as* is a transitive relation.²¹ Since each patch in the series looks the same in respect of colour as its successor in the series, the first patch looks the same as the last. And it follows from that theory that if x looks the same as y in respect of R and x looks F then y looks F. Since the first patch looks the same as the last patch in respect of colour and the first patch looks red then, it follows, the last patch looks red. But it doesn't – it looks yellow.

²¹ I take it that in the description of the series of patches 'x looks the same as y' means that the way that x looks in respect of colour is identical to the way that y looks in respect of colour. See Section 2.6.1 for why this is a transitive relation between x and y.

But Graff upholds Principle 1. That is, she thinks that 'looks the same as' expresses the *matching* relation. She says:

When 'looks' is taken in its phenomenal sense, the *looks the same as* relation is the same as, or at least akin to, what Goodman and others have called the *matching* relation. (p. 910)

As I will use it, 'matching' and 'looks the same as' stand for the same relation... (p. 911)

So what she is asking us to consider is a series of colour patches for which the first looks red and the last looks yellow, while each patch *matches* its successor in respect of colour, and what she is denying is the logical possibility of *this* series, while upholding that if any two patches *match* in respect of colour then if one looks red then so does the other.

Suppose, against Graff, that such a series of colour patches is indeed possible. That is, suppose that there are colour patches x_1 to x_{30} and an observer S such that, at the one time t , x_1 looks red to S , x_{30} looks yellow to S , while for each $k \in \{1, 2, \dots, 29\}$ x_k matches x_{k+1} in respect of colour for S . Moreover, suppose that S , by looking at the series of patches at time t , comes to believe that x_1 looks red (to her), that x_{30} looks yellow, and that for each $k \in \{1, 2, \dots, 29\}$ x_k matches x_{k+1} in respect of colour. S believes that x_k matches x_{k+1} in respect of colour. In many circumstances in which S believes that one object matches another in respect of colour for her, it does no harm for her to take it that the first *looks the same as* the second in respect of colour. If S is unaware of any paradox that might result from doing so, then she may well think that it will do no harm in this case, either, to take it that x_k *looks the same as* x_{k+1} in respect of colour. So she will think it harmless to take it that x_1 looks red, x_{30} looks yellow, while for each $k \in \{1, 2, \dots, 29\}$ x_k looks the same as x_{k+1} in respect of colour, and she may well say that this is the case. But suppose that S makes these claims in discussion with a logician, who points out to S that she has made inconsistent claims by presenting a sorites-style argument. Then S will learn, to her surprise, that in this context it *does* do harm to take it that for each $k \in \{1, 2, \dots, 29\}$ x_k looks the same as x_{k+1} in respect of colour. If she wants to do no harm in this context, then she will have to allow that for at least one $k \in \{1, 2, \dots, 29\}$ x_k does *not* look the same as x_{k+1} in respect of colour, even though it

matches it. So she will have to allow that for at least one $k \in \{1, 2, \dots, 29\}$ there is a colour C such that x_k looks C to her but x_{k+1} does not, even though she cannot come to know that there is, just by looking at x_k and x_{k+1} . When S looks at the series of patches, she acquires the knowledge that the first patch looks red, that the last patch looks yellow, and that the way that each patch x_k looks in respect of colour is so similar to the way that its successor x_{k+1} looks in respect of colour that she cannot come to know, just by looking at x_k and x_{k+1} , that there is an F such that x_k looks F in respect of colour but x_{k+1} does not, which is similar enough for her to think that it will do no harm to take it that x_k looks the same as x_{k+1} in respect of colour. What the sorites argument shows, to her surprise, is that there must be some k and some F such that x_k looks F but x_{k+1} does not, even though she cannot come to know which k that is, and that there is a context in which taking it to be otherwise does harm.

This response to the *phenomenal* sorites argument (Argument B) is very much like one we might give to the *non-phenomenal* sorites argument (Argument A). Suppose that S is told about a series of 4801 men m_1 to m_{4801} such that m_1 is tall (he is seven-foot tall), m_{4801} is short (he is three-foot tall), and for each $k \in \{1, 2, \dots, 4800\}$ m_k is one-hundredth of an inch taller than m_{k+1} . S believes that m_k is one-hundredth of an inch taller than m_{k+1} . In many circumstances in which S believes that two men differ in height by just one-hundredth of an inch, it does no harm for her to take it that the first *talls the same as* the second, by which I mean that either they are both tall or they are both not tall (i.e. that it is not the case that one is tall and the other is not). If S is unaware of any paradox that might result from doing so, then she may well think that it will do no harm in this case, either, to take it that m_k *talls the same as* m_{k+1} in respect of height. So she will think it harmless to take it that m_1 is tall, m_{4801} is short, while for each $k \in \{1, 2, \dots, 4800\}$ m_k *talls the same as* m_{k+1} , and she may well say that this is the case. But suppose that S makes these claims in discussion with a logician, who points out to S that she has made inconsistent claims by presenting a sorites-style argument. Then S will learn, to her surprise, that in this context it *does* do harm to take it that for each $k \in \{1, 2, \dots, 4800\}$ m_k *talls the same as* m_{k+1} . If she wants to do no harm in this context, then she will have to allow that for at least one $k \in \{1, 2, \dots, 4800\}$ m_k does *not* *tall the same as* m_{k+1} , even though they differ in height by just one-hundredth of an inch. So she will have to allow that for at least one $k \in \{1, 2, \dots, 4800\}$ m_k is tall but m_{k+1} is not,

even though she cannot come to know that that is so, just by knowing the heights of m_k and m_{k+1} . When S is told about the series of men, she acquires the knowledge that the first man is tall, that the last man is short, and that the height of each man m_k is so similar to the height of his successor m_{k+1} that she cannot come to know that m_k is tall but m_{k+1} is not, which is similar enough for her to think that it will do no harm to take it that m_k tall the same as m_{k+1} . What the sorites argument shows, to her surprise, is that there must be some k such that m_k is tall but m_{k+1} is not, even though she cannot come to know which k that is, and that there is a context in which taking it to be otherwise does harm.

These are both *epistemic* responses to the sorites arguments. In each case, S knows that adjacent items in the series are very similar in a certain respect (the way they look, ‘the way they tall’), but she does not know (in fact, *cannot* come to know, in the relevant kind of way) that they are the same in that respect (look the same, ‘tall the same’). She *thinks* she knows they are the same, because it has never done her any harm to do so: she has not experienced, and perhaps cannot imagine, a context in which items are so similar and yet cannot be taken to be the same. What the sorites arguments show her is that there *are* such contexts, and that she does not have the knowledge that she thinks she has. She does not have this knowledge in the non-phenomenal case because of her limited conceptual powers, and in the phenomenal case because of her limited perceptual powers.²²

²² See Williamson (1994) for an account of the epistemic view of the non-phenomenal case, especially Chapter 8 for why it arises from our limited conceptual powers.

Bibliography

- Armstrong, D. M. (1968), *A Materialist Theory of the Mind* (London: Routledge and Kegan Paul).
- Austin, J. L. (1962), *Sense and Sensibilia* (Oxford: Oxford University Press).
- Burgess, J. A. (1990), 'Phenomenal Qualities and the Nontransitivity of Matching', *Australasian Journal of Philosophy* **68**, pp. 206-20.
- Chisholm, R.M. (1957), *Perceiving: A Philosophical Study* (Ithaca, NY: Cornell University Press).
- Chisholm, R. M. (1965), "'Appear', 'Take', and 'Evident'", in R. Swartz (ed.), *Perceiving, Sensing and Knowing* (New York: Anchor Books), pp. 473-85.
- Clark, A. (1993), *Sensory Qualities* (Oxford: Oxford University Press).
- Crane, T. (1992), 'The Nonconceptual Content of Experience', in T. Crane (ed.) *The Contents of Experience* (Cambridge: Cambridge University Press).
- Davidson, D. (1967), 'The Logical Form of Action Sentences', in his *Essays on Actions and Events* (Oxford: Oxford University Press, 1980).
- Dummett, M. (1975), 'Wang's Paradox', *Synthese* **30**, pp. 301-24.
- Dummett, M. (1979), 'Common Sense and Physics', in G. F. Macdonald (ed.), *Perception and Identity* (London: Macmillan).
- Everett, A. (1996), 'Qualia and Vagueness', *Synthese* **106**, pp. 205-26.
- Gibson, Q. (1966), 'Is There a Problem About Appearances?', *Philosophical Quarterly* **16**, pp. 319-28.
- Goodman, N. (1951), *The Structure of Appearance*, 1st edn. (Cambridge, MA: Harvard University Press).
- Graff, D. (2001), 'Phenomenal Continua and the Sorites', *Mind* **110**, pp. 905-35.
- Hardin, C. L. (1988), 'Phenomenal Color and Sorites', *Nous* **22**, pp. 213-234.
- Jackson, F. (1977), *Perception: A Representative Theory* (Cambridge: Cambridge University Press).
- Jackson, F. and Pinkerton, R. J. (1973), 'On an Argument Against Sensory Items', *Mind* **82**, pp. 269-72.

- Knox, J., Jr. (1966), 'On Mr Nelson's Rejection of Sense-Data', *Ratio* **8**, pp. 94-5.
- Knox, J., Jr. (1967), 'The Logic of Appearing', *Inquiry* **10**, pp. 245-50.
- Knox, J., Jr. (1970a), 'Do Appearances Exist?', *APQ Monograph* **4**, pp. 79-101.
- Knox, J., Jr. (1970b), 'Speaking of Appearances', *Personalist* **51**, pp. 387-92.
- Lewis, D. K. (1979), 'Scorekeeping in a Language Game', *Journal of Philosophical Logic* **8**, pp. 339-59.
- Linsky, B. (1984), 'Phenomenal Qualities and the Identity of Indistinguishables', *Synthese* **59**, pp. 363-80.
- Mates, B. (1967), "Sense Data", *Inquiry* **10**, pp. 225-44.
- Nelson, J. O. (1964), 'On the Impossibility of Sense-Data', *Ratio* **6**, pp. 144-6.
- Paul, G. A. (1951), 'Is There a Problem about Sense-Data?', in A. Flew (ed.), *Essays in Logic and Language*, First Series (Oxford).
- Peacocke, C. (1981), 'Are Vague Predicates Incoherent?', *Synthese* **46**, pp. 121-41.
- Peacocke, C. (1993), *A Study of Concepts* (Cambridge, MA: MIT Press).
- Preyer, G. and Peter, G., eds. (2002), *Logical Form and Language* (Oxford: Oxford University Press).
- Price, H. H. (1964), 'Appearing and Appearances', *American Philosophical Quarterly* **1**, pp. 3-19.
- Raffman, D. (1994), 'Vagueness Without Paradox', *The Philosophical Review* **103**, pp. 41-74.
- Raffman, D. (2000), 'Is Perceptual Indiscriminability Nontransitive?', *Philosophical Topics* **28**, pp. 153-75.
- Sanford, D. H. (1981), 'Illusions and Sense-Data', in P. A. French, T. E. Uehling and H. K. Wettstein (eds.), *Midwest Studies in Philosophy* **VI** (Minneapolis: University of Minnesota Press), pp. 371-85.
- Soames, S. (1999), *Understanding Truth* (New York: Oxford University Press).
- Stanley, J. and Williamson, T. (2001), 'Knowing How', *Journal of Philosophy* **98**, pp. 411-44.
- Tappenden, J. (1993), 'The Liar and Sorites Paradoxes: Toward a Unified Treatment', *Journal of Philosophy* **90**, pp. 551-77.
- Travis, C. (1985), 'Vagueness, Observation and Sorites', *Mind* **94**, pp. 345-66.
- Williamson, T. (1990), *Identity and Discrimination* (Oxford: Basil Blackwell).

- Williamson, T. (1994), *Vagueness* (London: Routledge).
- Williamson, T. (2000), *Knowledge and its Limits* (Oxford: Oxford University Press).
- Wright, C. (1975), 'On the Coherence of Vague Predicates', *Synthese* **30**, pp. 325-65.
- Wright, C. (1976), 'Language-mastery and the Sorites paradox', in G. Evans and J. McDowell (eds.), *Truth and Meaning* (Oxford: Oxford University Press).