

Looks Unhelpful

William E. S. McNeill
University of Southampton, UK
will.mcneill@southampton.ac.uk

forthcoming in Mind

ABSTRACT¹

By looking at it you come to know that a thing is an apple. How? A natural answer is that this is down to how it looks – its superficial visual appearance. *Looks Views* treat our acquaintance with such looks as accounting for how visual knowledge is secured.

Here I argue that for many pairings of properties and perceivers *Looks Views* will turn out not to work. We can visually track many properties through huge variation in things' visual appearances. For such properties no kind of look will play the explanatory role *Looks Views* demanded.

Sometimes, then, you might secure visual knowledge that something is an apple but not by way of its superficial visual appearance. That conclusion may seem difficult to swallow. I argue that at least many of our intuitions about the roles of looks can nonetheless be assuaged.

property perception | perception of high-level features | epistemology of perception

1. Looks and Looking

Lots of our knowledge about the world is secured by looking. Consider a simple case:

REAL APPLE

There's an apple on the table and it's bathed in good light. It's a patchy red and green, ovally-round and slightly freckled. There's a thin brown stalk protruding from a small steep conical depression in its top.

¹ Specific thanks to the two anonymous reviewers for this journal and my colleagues Denis McManus and Daniel Whiting. Their constructive and insightful comments on earlier drafts pushed me to think much more carefully about the issues raised here, and develop these arguments in ways I could not have managed alone. Ancestors of this paper have had outings in various places, including the Joint Sessions in Durham and a Southampton perception workshop co-organised with Kurt Sylvan. With thanks to Kurt, Matt Soteriou, Alan Millar and Charles Travis for their thoughts and encouragement. And with thanks too to Alice – both for her patience and for persuading me that we cannot in fact see anyone's need of a good lawyer.

1 Call this apple *Real*.

2 On looking at *Real* you come visually to know that it is an apple.

3 Looking here is an activity – something you do. The knowledge you secure is a spontaneous
4 result of this activity. Now suppose I ask you how you knew that the thing was an apple.
5 Depending on the timbre of my question and the surroundings various answers might
6 suggest themselves.² A natural response – the one that interests me here – is for you to
7 say that *it looked like an apple*. What should we make of this answer?

8 You might simply be claiming that you’ve secured visual knowledge as opposed for
9 example to haptic, gustatory or testimonial knowledge. It wasn’t by its feel or taste or
10 because you’d accepted someone’s word for it that you knew. It was just by looking.³

11 Then again you might be trying to point out something of your visual reaction to the thing.
12 It looked *to you* to be an apple. It seemed to you an apple when you looked at it; it struck
13 you as one.⁴ You’re pointing to a specific, spontaneous psychological effect of looking at
14 the thing in the way you did.

15 However it can seem plausible that you’re instead trying to point me towards some relevant
16 features of the apple itself. It’s a patchy red and green after all; it’s slightly freckled with a
17 stalk. It’s a certain size and shape. These are aspects of the apple’s superficial visual
18 appearance – the *apple’s* looks. They are the kinds of features you might come across by
19 looking. They are the kinds of features that might cause things to strike you as they did.
20 And they are the kinds of features that I too might become acquainted with if I look.

21 Moreover there is an intimate connection between these features and the thing’s being an
22 apple. If it had been a tin of beans it would have looked very different. It would not have
23 struck you as an apple when you looked. And being a tin of beans, for example, does “seem
24 as though it fixes an object’s size, shape, and colour” (Parrott 2017: 1039).

25 A *Looks View* is any theory on which your acquaintance with objects’ superficial visual
26 features – their looks – accounts for the knowledge you secure just by looking. But

² ‘I work at a greengrocer’s’, ‘I love apples’, ‘I can see it’ (cf. Cassam 2007), ‘there’s nothing wrong with these eyes!’ and so on. Compare Austin’s (1961) example of the bittern.

³ For further discussion of *how you know* see Cassam 2007; French 2014.

⁴ See e.g. Tucker 2013, Moretti 2015 for overviews on phenomenal conservatism and seemings.

1 plausibly objects' looks also occasionally constitute the visual dead-ends which make our
 2 mistakes intelligible. How could I have known it was cake? It looked *exactly* like an apple.⁵
 3 Looks Views should be able to account for such failures. The cake only looks like an apple
 4 if it has the superficial appearance of an apple. And then it was its look that misled me.

5 In the next section I sharpen the commitments of Looks Views and set out various reasons
 6 why they should seem highly attractive. I will argue however that objects' looks often can't
 7 do the intended work – at least for many common features and where suitably sophisticated
 8 perceivers are concerned.

9 2. Looks Views

10 Primarily, Looks Views as here understood are normative. Concerning knowledge the
 11 thought is that a subject S visually knows that o is F only if:⁶

12 (i) o has a look L

13 (ii) o 's having L is appropriately associated with its being F

14 and

15 (iii) S is visually acquainted with o 's L .⁷

16 What it takes for a look to be appropriately associated with being an apple will require
 17 further discussion (see §3). But the basic idea is simple. If looking this way really is closely
 18 linked with being an F then taking the thing to be an F on the basis of its look is a *good idea*.

19 The question of how exactly to individuate the needed looks L is the primary concern of
 20 this paper. For now though we can identify an object's *superficial* visual appearances by
 21 employing the Visual Duplicate Test (Martin 2010: 200ff).

22 Take an apple and suppose that a master forger crafts a brilliant replica. Imagine the replica
 23 so good that no amount of careful visual inspection – no manipulation of lighting or angle
 24 of view – allows you to discern any visual difference between them. The two objects are
 25 *visual duplicates*.

⁵ There is currently a hit US TV series that centres around precisely this game – 'Is it Cake?'.

⁶ Plausibly if we're interested instead in the reasonableness of our perceptual beliefs, we can weaken (ii).

⁷ Cf. Parrott 2017 pp1029ff. Parrott discusses various further conditions that would have to be added. Also, note that I set aside questions about secondary perceptual knowledge (Dretske 1969) here.

1 To identify their superficial or basic looks reflect on what the forger would have to
 2 reproduce. She needs perfectly to replicate the apple’s patterns of shape, colour and
 3 reflectance. If any element of these differs then you’ll be able to tell them apart by looking.
 4 Shape, colour and reflectance, then, are amongst things’ superficial visual appearances. Yet
 5 the forger would neither need to nor could they create an actual apple. So *being an apple* is
 6 not amongst the apple’s basic looks.

7 Individuated this way *looks* turn out just to be objects’ sensible qualities or observational
 8 properties.⁸ And such looks have what I will call *perceptual privilege* over other features of
 9 visible things. What the visual duplicate test uncovers are all and only those properties
 10 which interact directly with light. In this way objects’ looks will at least be causal
 11 intermediaries to visual knowledge of any of their other features.

12 Accepting the necessary conditions above on visual knowledge is compatible with stark
 13 epistemological disagreement. For Millar (2000; 2011) the needed machinery here is simply
 14 the cause of an immediate response to looking at the thing (2011: 338). The knowledge
 15 you secure is basic and non-inferential (2000: 77).⁹ By contrast McGrath (2017; 2018) treats
 16 *L* as your *evidence*. In conjunction with knowledge of the association between *L* and *F* this
 17 evidence puts you in a position to secure only mediate or inferential knowledge that *o* is *F*
 18 by looking (2018: 131).

19 In spite of these differences Millar and McGrath do also agree that our acquaintance with
 20 things’ looks allows us to make sense of a certain kind of mistake. After all even if having
 21 a look *L* is ‘a very highly reliable indicator’ (Millar 2011: 339) that it’s an *F* the association
 22 might break down in particular cases. Perhaps most things that are a patchy red and green,

⁸ Compare Peacocke’s observational concepts (1983 chp.4); McGrath’s perspective-invariant looks (2017, p36 and 2018, p111). See also Martin 2010, p203.

This gives us a very simple picture of what looks are. Brewer (2011) and Schellenberg (2008), for example, argue instead for a standpoint relative model (compare Millar 2000, p79 and McGrath 2017, p37; 2018, p130.). This too would move us away from a focus merely on how things strike you. And no doubt *Real* will look different from different vantage points. You can see the maggot hole but I can’t. Sometimes, such perspectival looks are going to play a crucial role.

However for reasons set out in §2.1 below a focus on these simple looks suffices here. The standpoint invariant look of an object is what the object itself is contributing to its look from this angle, in this lighting. It’s that input to the equation that is actually a feature of the object about which you come to know. Moreover it transpires that the needed looks would have to be both characteristic (§3.3) and common (§4.1) to *F*s – much more likely for standpoint invariant than for standpoint-relative looks.

⁹ Compare Shieber 2017.

1 slightly freckled, oval-round and so on are apples. The occasional forgery may well dupe
2 you.

3 For McGrath you'd have secured the same justification here as if the thing had in fact been
4 an apple. Millar disagrees. Nonetheless he accepts that there's an important sense in which
5 your belief would be 'reasonable' (2011: 345). And again it's the point of intersection that
6 interests me here – the idea that the looks we're acquainted with allow us to *make sense* of
7 our visual reactions.

8 2.1. Motivating Looks Views

9 One way to motivate Looks Views is to appeal to plain common sense. It's deeply plausible
10 that if you're to identify *F*s by looking, *F*s must look some way and you need to know how
11 *F*s look. Consider how absurd it might sound to deny this. Imagine the apple looked *no way*
12 *at all* to you. You couldn't have known anything about it by looking.¹⁰ Or suppose you
13 don't know what apples look like. Then surely you can't know *Real* is an apple just by
14 looking at it (McGrath 2017:8).

15 Looks Views can also be motivated by contrast with two alternatives. At first sight they
16 offer advantages both over views focusing solely on how things strike you and over views
17 where your acquaintance with objects' rich features is treated as *sui generis* or basic.

18 When we talk of how things look we tend to identify relatively stable features that can be
19 readily shared – the patchy red-and-green; the freckles. We can often find agreement as to
20 how things look even if our perspectives differ profoundly. I don't need to push you out
21 of your spot to agree that the avocado looks green, the tin shiny.

22 And intuitively we can also disagree. We can disagree as to whether the mango really looks
23 ripe; whether the dress looks at all blue (McGrath 2018: 119). But how things strike you –
24 how things look *to you* – is a personal matter. Those visual reactions are not things we can

¹⁰ A couple of caveats. Perhaps you don't really see silhouetted objects. If not they surely couldn't look any way to you. But you might tell that a silhouetted thing was an apple by looking. I leave these complications aside here. Secondly we can at least imagine a superblindsighter saying the apple looks no way to them but that nonetheless they know by looking that it's an apple (and thanks to an anonymous referee for this point). Given that they concern the structure of looks themselves or perspectives on them rather than the phenomenology of looking, the arguments against Looks Views offered here should work over such blindsight cases too.

1 share in or debate. So it's intuitive that we do have an interest in objects' looks as distinct
2 from our own psychological reactions to those objects (Millar 2000: 78).

3 At the same time we've noted that objects' basic looks are perceptually privileged. It's
4 common to treat only these properties as 'strictly visible' (Millar 2011: 339); to accept that
5 'perceptually manifest facts concern only the way things appear' (Millar 2000: 74)¹¹. And
6 even if we allow that facts about objects' rich features are perceptually manifest we might
7 still agree that 'all perceptual knowledge is a response to what is perceptually manifest in
8 this fairly strict sense' (Millar 2000: 77). Vision can't sidestep things' looks. Looks Views
9 simply respond to the thought that it must then be somehow *via* these looks that we secure
10 further knowledge.

11 These explanatory advantages seem manifest too if we reflect on two kinds of perceptual
12 mistake. Our master forger has secretly altered the appearance of a cake so that it strikes
13 everyone who sees it as an apple, even on close inspection. You look at this thing and it
14 strikes you too as an apple. Contrast a case in which, having unwittingly taken a
15 hallucinogen, an undoctored victoria sponge strikes you as an apple.

16 In both cases the cake strikes you as an apple. In both cases you believe falsely that it's an
17 apple. So the contrast cannot be explained in terms either of what you perceptually know
18 or how things strike you.

19 Yet in the first case your false belief seems reasonable or explicable in a way that it wouldn't
20 in the second. In the first but not the second case the cake really does look like an apple.
21 You've responded appropriately to how things look and your mistake is just bad luck.
22 However the locus of the problem in the second case is not how the thing looks but *your*
23 *brain*. It's pathological.

24 A final benefit of Looks Views is that they promise to give a plausible account of
25 perception's scope. We can secure perceptual knowledge of more than merely objects'
26 basic looks. We can tell by looking that the thing's a tin as well as that it's shiny. The Looks
27 View is that we can make sense of that where we find some intimate relationship between
28 perceptually privileged looks and the other features we come to know of.

¹¹ Compare Parrott 2017, p1024

1 But that also gives us the tools to withhold from an overly permissive “non-inferentialist”
 2 view of visual knowledge (Millar 2000: 74). Even if we’re happy to *say* that the book looks
 3 unhelpful or that the red-faced stranger looks to be in need of a good lawyer *being unhelpful*
 4 and *being in need of a good lawyer* aren’t intimately related to any particular look. In line with
 5 commonsense the Looks View implies that you’d not be securing visual knowledge in these
 6 cases.¹²

7 2.2. Plan

8 In spite of these prima facie advantages I will argue that Looks Views aren’t viable. We can
 9 sometimes recognise apples on sight but not *on the basis of* their look. To see what’s at stake
 10 here compare our visual acquaintance with things’ colours and shapes. If you’re colour
 11 sighted then on looking at the book you can tell that it’s red. How? A simple answer is that
 12 it’s because you saw its colour; because you’re visually acquainted with the very redness
 13 about which you know.

14 That answer isn’t trivial. Archivists in a really dingy vault might see that the book is red on
 15 seeing the luminous mould that only colonizes covers dyed red.¹³ With such poor lighting
 16 the redness itself doesn’t visually stand out to them. The archivists see that the book is red
 17 but only *by* visual acquaintance with the luminous mould.

18 The Looks View is that your visual knowledge of *Real* is analogous to the archivists’
 19 knowledge of the book’s colour.¹⁴ You see that the thing is an apple but only *by* seeing its
 20 superficial visual appearance. Just as awareness of the luminous mould accounts for the

¹² Perhaps where you spontaneously make such judgements you don’t secure knowledge but something less. Or perhaps you’ll secure knowledge just not perceptual knowledge. See §5 for further discussion.

¹³ This adapts an example in Martin 2010, p175. See also McNeill’s (2012, p587) example of the achromat’s ability to spot that the traffic light’s red.

¹⁴ There are some knotty metaphysical questions here – in particular concerning how to individuate the featural relata. As the archivist example shows visual acquaintance can’t just be to facts. Johnston (2006; 2018) argues that it is to states of affairs. However it’s not obvious that the archivists fail to see the red of the book. They’re disposed to treat the book as red on looking. In this sense it looks red to them.

A similar concern arises if we’re acquainted with property *instances* (see e.g. Ivanov 2017; Schellenberg 2018, p15). The archivist can see the red book. They can see that patch of red which is its cover. Visual acquaintance with such patches doesn’t distinguish your awareness of the book’s colour in good light from the archivist’s. If that’s all it took to see the red of the book, doing so wouldn’t be much use. For now I stick with the commonsense distinction. However exactly we are to model this, you see the book’s colour but the archivist does not. That acquaintance accounts for your knowledge in a way that it couldn’t account for the archivists’.

1 archivists' knowledge of the book's colour, seeing *Real's* look is supposed to account for
2 how you know that it is an apple.¹⁵

3 I argue here that such a View will fail to account for your visual knowledge of many rich
4 features. The problem gets going whenever you're good at spotting those features in real
5 time and across wide variations in how they are visually presented.

6 In their vault the archivist relies on a very specific set of circumstances to recognise redness.
7 They could easily lose that ability if the mould dies off or colonizes other pigments. But
8 out in the world the colour sighted can see redness across all sorts of ways things look –
9 with only slightly better lighting than the vault's, through fog or strobing shadow.

10 In general, knowledge *on the basis of* a look will be fragile. Remove, undercut or rebut that
11 evidence and there's no knowing. But just as you can spot redness in all sorts of lights your
12 capacity to recognise apples is *robust*. What emerges is a view on which you can sometimes
13 know that a thing's an apple not *by* its look but by acquaintance with that very feature of it
14 – it's being an apple.

15 Consider Austin's pig. If all I can see are some 'pig-like marks on the ground' (1962: 115),
16 I have evidence of a pig. If I know there's a pig this is because I'm acquainted with those
17 marks. However when the pig comes into view this 'doesn't provide me with more *evidence*
18 that it's a pig. I can now just see that it is, the question is settled' (*ibid.*).

19 How do I know there's a pig there? It's not *just* a matter of its visual presence. I might see
20 the pig without being able to recognise pigs.¹⁶ If in my ignorance someone sticks a sign on
21 its back that says 'pig' I may know that it's a pig but *still* on the basis of evidence. The
22 question – as it were – would not be settled.

¹⁵ I can't be precise about what it is for one thing to account for another. But by comparison if the fact that I strike the match accounts for its lighting it's because we can take for granted that the other conditions – enough oxygen, not too much wind and so on – are met. By contrast the mere fact of my throwing the die cannot account for its landing on six. Only a much more complex set of features – starting position, muscular forces and so on – could do that.

¹⁶ Travis (2004) for example argues that such recognitional abilities would have to be conceptual. While the arguments here are compatible with that view they don't imply it. Even if they don't have the associated concepts a horse might have a robust and fine-grained ability to spot apples; a lion to spot pigs.

1 It's only settled, I argue, if I'm a reliable spotter of pigs – if I have a general capacity to
 2 recognise pigs when I see them. And if seeing the pig settles the question of whether it's a
 3 pig for me then I'm visually acquainted with its being a pig; with its pigginess.

4 The plan for what follows is this. I focus first on how to understand condition (ii). What
 5 link must hold between having a look L and being F , if acquaintance with the former can
 6 account for knowledge of the latter? It soon transpires that L cannot be o 's specific,
 7 determinate basic look. What we need is a better understanding of what L is.

8 I argue in §3 that an appropriate look L will need to account not only for how by looking
 9 you know or why you're warranted in judging that o is F . It will have to account also for
 10 why the object actually strikes you as an F .

11 In §4 I argue that where sophisticated apple spotters are concerned no such look exists.
 12 The problem emerges when:

- 13 a) your disposition to judge on looking that things are F is *robust* over very different
 14 patterns of basic look (it's 'stimulus invariant') and
- 15 b) your dispositions are *fine-grained* – there are objects which look similar to those
 16 you're disposed to judge on looking are F s which you're *not* disposed to judge
 17 on looking are themselves F s.¹⁷

18 Where your dispositions are this complex, acquaintance with no look L could account for
 19 them.¹⁸ If I'm right that only looks that play this causal role could play the normative role
 20 that Looks Views invoke, Looks Views aren't viable.

21 The argument relies on descriptive claims as to the situations in which we're disposed to
 22 form visual apple beliefs – or not. A defence of Looks Views might push back on such
 23 claims (§5.1). The danger is that we lose sight of a natural understanding of the roles
 24 perception plays. It's deeply plausible that visual perception can deliver knowledge in the

¹⁷ The distinction here appears in McNeill 2015, p1436.

¹⁸ Another way to run this would be to think in terms of similarities to a prototypical apple or set of them (cf. Brewer 2011, especially chp5). The arguments here should translate. Things which are visibly apples may be very dissimilar in many respects. Things which are visibly *not* apples will bear countless visual similarities to apples (cf. Travis 2004). But if all we mean is that these things similarly strike *you* as apples, any objective similarity to prototypical apples ceases to pull weight. We'd have retreated from how things themselves look to how you react when looking at them.

1 wild and on the fly; across huge variations to the light array. As long as perception is robust
2 and fine-grained in these ways Looks Views won't work.

3 Yet there's something troubling about all this. After all I've outlined how intuitive Looks
4 Views sound; how absurd it can feel to deny them. The final sections aim to soothe the
5 nerves. We can account for many of those intuitions without invoking *looks* (§6). It helps
6 too to reflect on an analogy with less controversial and more familiar claims about
7 perceptual constancy (§5.2).

8 3. Distinctive and Characteristic Looks

9 Primarily, Looks Views are normative. They aim to account not for why you *form* the visual
10 beliefs you do but only why – when they are – those beliefs are reasonable; how they might
11 secure you knowledge. I argue here though that in order to fulfil that normative brief the
12 relevant Looks in fact need to account for the spread of your visual dispositions.

13 Looks Views gave us two normative conditions on securing visual knowledge that *o* is *F*:

14 (i) *o* has a look *L*

15 and

16 (ii) *o*'s having *L* is appropriately associated with its being *F*

17 Both conditions need unpacking. I begin with (ii).

18 3.1. Distinctiveness

19 We've seen that if on acquaintance with *o*'s look *L* you come to know that *o* is *F*, *L* must
20 be an indicator of a highly reliable kind to the presence of *F*s. In other words the relevant
21 looks must be

22 *distinctive* a look *L* is distinctive of *F*s just if, in the circumstances, most things
23 that have *L* are *F*¹⁹

24 Only if there's some tight connection between *L* and *F* could acquaintance with *L* make it
25 a good idea to treat *o* as *F*. But the connection can only be so tight. If the constraint were
26 that *only* *F*s have *L* the mere possibility of visual doppelgangers would threaten visual

¹⁹ Cf. Parrott 2017, p1025; Millar 2011, p339.

1 knowledge of *F*s. Yet even though there could be doppelgangers for apples we can
 2 sometimes know by looking that things are apples. So it must be possible for things which
 3 are not *F* to manifest *L* too.

4 Looks Views deal with this by making distinctiveness contextual. Consider

5 **SCHMAPPLES**²⁰

6 Schmapples are poisonous fruits very different from apples both genetically and
 7 in their uses. Schmapples are visually indistinguishable from apples. But in
 8 schmapple country they're common and turn up in the same kinds of places as
 9 apples.

10 Suppose yourself unwittingly transported there. You'd not be *safe*. You'd be
 11 liable to mistake schmapples for apples and the results could be fatal.

12 In this situation you might well form true beliefs that things are apples by looking. You
 13 wouldn't thereby count as knowing that these things were apples. There are plenty of
 14 schmapples around and any one of them would fool you. Apple doppelgangers – hence
 15 false beliefs – are too close at hand.

16 Looks Views give us a way to account for the problem in SCHMAPPLES. We can say that
 17 apples and schmapples look the same – that is, *have the same look* (Parrott 2017: 1027). In
 18 schmapple country this look – being shared by apples and schmapples – cannot be
 19 distinctive of either. It misses condition (ii). Yet around here very few things have the same
 20 looks as apples. Around here that look meets (ii). And around here we can secure
 21 knowledge of apples by looking.

22 However this pushes us to ask what the look *L* is that apples and schmapples share – to
 23 focus on condition (i). After all schmapples, like apples, all look different from each other.
 24 All these fruits have had a somewhat chaotic, outdoor upbringing. They'll be slightly
 25 different shapes, have different patterns of colour, freckles, bumps, scabs and rust.

²⁰ From Parrott 2017, p1027. Compare McGrath 2017, p38; Millar 2011, p343.

3.2. The Safety Test

What we need is a way to test whether candidates for L really play the needed role. We can fashion such a test by holding fixed whether you know by looking that o is F in some context and determining whether L is distinctive there.

For suppose a Looks View is correct and on acquaintance with L you form the true belief that o is F . Then all else being equal you know that o is F if and only if L is distinctive of F s in this context. Running this engine in reverse, a candidate look L fails what I call the *Safety Test* if on acquaintance with L and all else being equal either:

ST1: you visually know that o is F yet L is not distinctive of F s in this context.

ST2: you only truly believe but don't know that o is F , yet L is distinctive of F s in this context.

If either condition is met then acquaintance with L could not be what accounts for whether you secure visual knowledge.

To see the test in action consider first a natural candidate for L . You get a really good look at an apple, in good light. Call the highly detailed and determinate look with which you're thereby acquainted L_D . While repeatable in theory, in practice L_D is unique. Nature won't produce a visual duplicate and the forger's fee would be far too high. So everything that has L_D is an apple. But then even if you are in schmapples country L_D would be distinctive of apples – this one in particular.

However in SCHMAPPLES you'd only truly believe that the thing with this look was an apple. You wouldn't know it. So via ST2 L_D fails the Safety Test. The reason that L_D delivers the wrong verdict here is that acquaintance with L_D wouldn't reflect how lucky you were to have formed a true belief in SCHMAPPLES.

3.3. Characteristicness

A first try at correcting the problem – of identifying the needed look – would be to specify that it be not only distinctive but also

characteristic : L is characteristic of F s just if in the circumstances most F s have L ²¹

²¹ Compare Millar 2000, p78; McGrath 2017, p4

1 This condition excludes L_D – only one apple had that look. And it captures something of
 2 the intuition that visibly identifiable apples can't look just *any* way. If there were no
 3 'regularity to the way they look' (Smith 2015: 2) we'd not be able to spot them just by
 4 looking.²²

5 However there are many looks both characteristic and potentially distinctive of apples that
 6 would nonetheless fail the Safety Test. Take the look L_{RR} – *roughly palm-sized, reddish and*
 7 *roundish*. Most apples have this look – it is characteristic of apples. And it's just a
 8 determinable of most apples' more determinate looks. So mostly when you look at an apple
 9 you'll be acquainted with L_{RR} . Moreover it's easy to imagine situations where L_{RR} would be
 10 distinctive of apples.

11 Yet by ST1 L_{RR} fails the Safety Test. In the vicinity of ripe plums you can still perfectly well
 12 know by looking that things are apples. But ripe plums share L_{RR} with apples. So in the
 13 vicinity of ripe plums L_{RR} would no longer be distinctive of apples. L_{RR} delivers the wrong
 14 verdict because acquaintance with L_{RR} doesn't reflect how unlucky you'd have been to form
 15 false beliefs around here.

16 The Safety Test pushes us to discover what kinds of things actually count as doppelgangers
 17 for you. A doppelganger in the relevant sense is something that strikes you as an F – that
 18 you're thus disposed to treat as an F – even though it's not.²³ SCHMAPPLES shows that
 19 doppelgangers need not be visual duplicates. In fact doppelgangers that are visual
 20 duplicates will be incredibly rare, at least for natural objects like apples.

21 L_D predicts that nothing would count as a doppelganger – that's not right. But L_{RR} predicts
 22 that too many things would count as doppelgangers – anything reddish, roundish and
 23 roughly palm-sized; things you can perfectly well tell by looking are not apples; things that
 24 don't strike you as apples when you look.

25 To work out what your apple doppelgangers are is to ascertain the pattern of what actually
 26 strikes you visually as an apple, in which situations. And it's this which determines whether

²² As a causal matter a characteristic look is something you'd have a hope of both learning to associate with apples and of spotting on many different encounters with apples. I come back to this in §4.1 below.

²³ In line with the motivations for Looks Views and the role looks are taken to play we can exclude pathological cases; should restrict ourselves here to situations in which your perceptual machinery is working well.

1 or not your true beliefs round here are lucky; how likely you'd be to have believed falsely
2 by looking.

3 In sum Looks Views can answer the question of how you know what you do by looking
4 only if they identify an appropriate look (from (ii)). But appropriate looks must pass the
5 Safety Test. And whether a particular look passes the Safety Test is an empirical matter. It
6 has to do with the spread of occasions when on looking things actually strike you – rightly
7 or wrongly – as apples.

8 4. Spotting Apples

9 If visual acquaintance with L accounts for its striking you as an F you must somehow
10 associate having L with being F .²⁴ Here I tease out the conditions that would need to be
11 met if this association, alongside visual acquaintance with L , could account for which things
12 strike you as apples when you look.

13 4.1. Common and Distinguishing Looks

14 Consider first the causal correlate of *characteristicness*:

15 *common* a look L is common to F s for S just if only those things with L strike
16 S as F s

17 Lots of things without L_D strike you visually as apples – all the schmapples when you land
18 in schmapples country, for example. *Common* shows us how we can dismiss L_D as the
19 relevant look.

20 An initial problem for *common* is that you can spot all sorts of apples by looking. To take
21 some simple examples, not all will be a patchy red and green. The Granny Smiths are a tall,
22 shiny plain green. The Egremonts are a squat, uniform, dull yellowy brown. Given how
23 different these varieties are in their looks there'll be no coherent look they'd have in
24 common that isn't incredibly vague. It would have to be a look so determinable they'd have

²⁴ For Millar and McGrath the way you realise this association will be very different. For Millar it's merely an enabling condition for an immediate visual response – perhaps the result of a sub-personal visual mechanism. For McGrath it will be an indispensable aspect of the reason-giving structure that makes your perceptual knowledge mediate.

1 it in common with countless other roughly palm-sized objects – with many visibly
2 identifiable potatoes for example. Looks this indeterminate aren't distinctive of apples.

3 But it's also just not plausible that there *need* be a single look which all these things share,
4 in virtue of which they all strike you as apples. Learning to identify not only Braeburns but
5 also Granny Smiths and Egremonts as apples may well involve associating a number of
6 different looks with *being an apple*.²⁵

7 Crucially though there will have to be some limits. The number of different looks will have
8 to be learnably finite. For each such look you (or your visual system) will have to learn to
9 associate that look with being an apple. That task must be accomplished – hence
10 accomplishable – with finite cognitive resources. We couldn't for example just gesture at
11 all the different determinate looks of everything that has struck or will strike you as an
12 apple. You couldn't have managed to associate all these indefinite numbers of distinct looks
13 with being an apple.

14 Where it's plausible that we've learnt to associate several distinct looks with being *F* we can
15 treat *L* as a finite, learnable disjunction of discrete, common looks. But while *common*,
16 suitably hedged, seems to give us a plausible start in understanding how things strike us it
17 would need bolstering.

18 L_{RR} is a look common to apples ... but also to visibly identifiable plums, cricket balls and
19 so on. *Common* doesn't capture the thought that if acquaintance with *L* accounts for *o*'s
20 striking us as an *F* then acquaintance with *L* must dispose us to judge the thing an *F*. We
21 need –

22 *distinguishing* *L* is distinguishing of *F*s for *S* just if on acquaintance with *L* and all
23 else being equal *o* will strike *S* as an *F*

24 This condition would allow us to set aside L_{RR} . Many things which are visibly reddish,
25 roundish and roughly palm-sized will not in fact strike you as apples. It's that empirical fact
26 that reassures us that the proximity of plums doesn't stop you knowing on looking that a

²⁵ Compare Martin's example of the *taste of French cheese* (2010, p177). While Camembert, Roquefort and Brie taste very different from each other they might all still come to strike you as *French* when you taste them. Tasting French, here, is an association of a disjunction of very different but individually common and distinguishing tastes. See also McGrath 2017, p37.

1 thing’s an apple. Looking doesn’t dispose you to the false belief that plums are apples. So
 2 the presence of plums doesn’t have a significant impact on how lucky you are to form true
 3 apple beliefs by looking.

4 The discussion so far motivates the search for a look L which is not only characteristic and
 5 distinctive of apples but also common to and distinguishing of apples for you. I argue next
 6 that there is no such look – at least where sophisticated apple spotters are concerned.

7 4.2. Your Robust and Fine-Grained Capacity for Apple Spotting

8 At first sight it might seem that unlike L_D or L_{RR} the following look meets these conditions
 9 –

10 L_{CA} a patchy red-and-green, slightly freckled and ovaly-round with a steep
 11 conical indent at the slightly wider end

12 L_{CA} is a classic, archetypal way for an apple to look. It’s characteristic of them – or at least
 13 of *ripe Braeburns*. It’s also likely to be distinctive. You wouldn’t catch an avocado looking
 14 this way. Even so it will turn out that acquaintance with L_{CA} won’t do the required causal
 15 work. As billed L_{CA} isn’t *distinguishing* of apples for you. Consider

16 VISIBLE FAKE

17 There is an object on the table. It is oval-round and its surface is a patchwork of
 18 red and green. It is shiny and slightly freckled. It has a thin brown stick protruding
 19 from its indented top. It is not an apple. It is a mildly well-crafted model.

20 Call this object *Fake*.

21 Of course some fakes – indefinitely many – would fool you. But this one doesn’t.

22 It looks fake to you. On seeing *Fake* you know that it is a fake and not a real apple.

23 Sometimes when a thing strikes you as being a fake apple you might be able to point to
 24 helpful contextual clues. You’re in a craft shop. You saw the same thing a month ago and
 25 it hasn’t rotted. It’s labelled “£20” or “wooden reproduction”. But this isn’t always how
 26 things go. Visual knowledge that something is a fake apple can be what *informs* you that
 27 you’re in the craft section, that it won’t rot, that it will cost too much.

1 *Fake* and *Real* are different in their look, and obviously so. But then again none of the
 2 visibly identifiable apples are visual duplicates either. Almost everything that strikes you as
 3 an apple is visually unique.²⁶ The fact that *Fake* and *Real* visibly differ does not by itself
 4 account for the different ways they look to you to be.

5 And *Fake* really does share with *Real* the determinable look L_{CA} . *Fake* really is apple-like in
 6 shape and size. It really is a patchy red and green, slightly freckled and so on. It's just that
 7 it is also obviously, visibly *not an apple*. Given that the two objects strike you differently but
 8 share L_{CA} , L_{CA} is not distinguishing of apples for you.

9 And so by ST1 L_{CA} fails the Safety Test. *Fake* is not an apple doppelganger for you. You
 10 can still secure visual knowledge that things are apples even with *Fake* around. Yet if *Fake*'s
 11 around, L_{CA} is no longer distinctive of apples. Having the look L_{CA} is not appropriately
 12 related to being an apple.

13 Two promising responses jump out though. We might give up on L_{CA} and search for a
 14 look which is distinguishing of those things that do actually strike us as apples – a *real* apple
 15 look L_{RA} . Alternatively we might focus on what's special about *Fake*.

16 Start by trying to individuate L_{RA} . This must still be a look shared by all those ripe
 17 Braeburns that strike you as apples. It must *also* be something that visibly fake apples lack.
 18 It's difficult to imagine what this look could consist in.

19 The problem is that your visual reactions are highly *robust* – ‘stimulus invariant’. Even ripe
 20 Braeburns can come in all sorts of shapes, sizes, colours and sheens. Braeburns that differ
 21 in all these respects can nonetheless all visibly be apples. So we can't collar any *very*
 22 particular shape; any *very* specific pattern of red and green patches or distribution of
 23 freckles; anything much more determinate than L_{CA} . Yet the patterns of sensible quality
 24 those Braeburns that are visibly apples all have in common with each other – the generic
 25 and permissive *oval-round* and so on – are things they could surely have in common with
 26 *Fake*.

²⁶ Imagine you came across a pair of perfect visual duplicates. Even if they both strike you as apples when you look you have here good reason not to hold the beliefs you're thereby disposed to hold. At least one of them is surely a fake.

1 And if it's hard to imagine what look could account for even saleable Braeburns striking
 2 you as apples consider the odds of finding a look that distinguishes *mouldy* apples for you.
 3 Every apple rots in its own special way. Mangled mouldy lumps that strike you visually as
 4 apples are visually *very* diverse.

5 Given all this it might appear more promising to focus on what's odd about *Fake*. Can you
 6 just make out a seam? Is that the grain of the wood showing faintly through? Perhaps it
 7 looks just a bit too shiny or symmetrical.²⁷

8 Fakes are certainly unusual. They are individuated in part by their intended superficial
 9 similarity to apples. Given this perhaps the slight seam or the thing's shininess are acting
 10 as visual *defeaters* of L_{CA} 's natural import. It looks *sort of* like an apple – just not quite.²⁸

11 But on what grounds could such visual details really behave as defeaters? For an aspect of
 12 a thing's look to defeat the link between L_{CA} and its being an apple it must somehow *rub*
 13 *up* against L_{CA} . But as currently individuated L_{CA} is silent on seams and symmetry; is
 14 compatible with wide variations in shininess. To be defeated by *Fake*'s look L_{CA} itself will
 15 require some serious enrichment.

16 We've noted an initial problem with that. Our dispositions are robust across all sorts of
 17 looks. Enriching L_{CA} threatens its characteristicness. A parallel problem here is that for
 18 fakes there is just such a variety of imaginable tells. Is it that this one is floating just a bit
 19 too high in the water? We (or our visual systems) were meant to be able to learn to associate
 20 L_{CA} with being an apple; to learn what apples (at least of the Braeburn variety) look like.
 21 We can't do that if L_{CA} is unlearnably complex.

22 A no-less fundamental problem is that properties like *lacking a seam* don't have perceptual
 23 privilege (p4). L_{CA} was supposed to characterise things' superficial visual appearances. But
 24 *being seamless* is not something that interacts even indirectly with the light array. It's no aspect
 25 of any object's overall look – nothing the perfect forger either could or would need to
 26 recreate.

²⁸ With thanks to an anonymous reviewer at this journal for framing things this way.

1 VISIBLE FAKE shows that our visual reactions are quite *fine-grained*. Relatively small
 2 differences in how two things look – those which distinguish *Real* and *Fake*, for example –
 3 can dispose them to strike us as objects of very different kinds. Together these two tenets
 4 of our dispositions – that they’re robust and fine-grained – tell against the view that
 5 acquaintance with any merely superficial look could account for how things strike us when
 6 we look.

7 The relevant look must be *common* to some variety of things that strike us as apples. The
 8 more varied are the things that strike us this way the more determinable that common look
 9 would have to be. But the relevant look must also be *distinguishing* of apples for us. The
 10 more fine-grained our dispositions to believe that things are apples the more detail we’d
 11 need to add; the less determinable the look could be.

12 4.3. Real-World Visual Knowledge

13 The argument so far has focused on situations where we’ve been given the chance to have
 14 a really good look at the thing. Even in such visually benign contexts we won’t always be
 15 able to account for the knowledge we secure in terms of how things look.

16 But the argument has also pushed us to reflect on our actual dispositions to form beliefs
 17 by looking. Plausibly we’re disposed to make visual judgements quickly and in the cut and
 18 thrust of fleeting and partial visual encounters with the world – when conditions are far
 19 less than benign.

20 A thing might strike you as an apple even through strobing shadow, in fog, at twilight, or
 21 while deep in the long grass. So imagine as against the arguments of the previous sections
 22 that we could find some independent means of identifying a look appropriately associated
 23 with *F*s. It’s far from clear we need to be acquainted with it on occasions where we’re
 24 disposed to believe the thing an apple by looking; on occasions where we know it is. The
 25 Safety Test makes (iii) a problem even if (ii) were not. Consider

26 BOBBING APPLE

27 There’s an unblemished ripe Braeburn caught spinning in the churning wash of
 28 a small stream.

29 Call this apple *Spin*.

1 On catching sight of it *Spin* strikes you as an apple. You come to know there's
2 an apple in the stream.

3 In some similar situation you may have been able to identify aspects of context that played
4 in to how things seemed – the stream's just under an apple tree; you'd heard an apple truck
5 had tipped over. But that's not how things unfolded on this occasion. Here the thing's
6 striking you as an apple is what prompts you to look up and around in search of the apple
7 tree; to wonder if there must have been a storm in the night to dislodge this apple so early
8 in the season.

9 *Spin* has a classic appley look. But from your vantage point you can't really make this out.
10 *Spin*'s definitely reddish and roundish but given the bubbles, the rush of water and the way
11 it's being turned over and round you'd be hard pushed to report any further detail.

12 That doesn't mean that any old reddish, roundish thing caught there would have struck
13 you as an apple though. The details matter. Would *Fake* have floated a little higher in the
14 water, spun faster, caught the light differently? For some reason it just didn't strike you
15 visually as a cricket ball or plum when you looked. In not dissimilar circumstances perhaps
16 it would have, or perhaps you'd have been less sure. But here you'd have been amazed to
17 discover it was anything but an apple.

18 *Spin* is designed to bring out just how robust your dispositions might be to believe things
19 are apples when you look. But if you can find yourself in such situations it really is going
20 to be impossible to identify a look that both passes the Safety Test *and* with which you're
21 acquainted when on looking things strike you as *F*s.

22 Together what VISIBLE FAKE and BOBBING APPLE bring out is that in one way or the
23 other the details matter. It's something quite specific to *Real*'s look that would mark it off
24 from *Fake* – something that sets *Real* apart from other visually identifiable apples too;
25 something uncommon and uncharacteristic. And it's something quite specific about *Spin*'s
26 context – not its basic look – that will causally account for its striking you as an apple. In
27 this context *Spin* itself can't manifest any distinctive or distinguishing features.

28 But the relevance of all the idiosyncrasies of these different cases means there's no stable
29 look *L* in terms of which we could work out what your apple doppelgangers are. *No* look

1 passes the Safety Test. And in turn that means there's no look L that we can say is
 2 appropriately associated with being F . No look can account for how you know things are
 3 apples by looking – at least when your dispositions are robust and fine-grained. Looks
 4 Views aren't viable.

5 Yet notice that this isn't because how things happen to strike us is *random*. It's because
 6 we're really good at spotting apples. It's rare that anything else would strike us as an apple
 7 when we look and we can visually pick apples out in all sorts of contexts. Our visual
 8 dispositions manifest a fine-tuned capacity to spot these rich features of our environment,
 9 in real time, in the mess and clutter of the real world.

10 5. Perception's Scope

11 The situation we're left with may still seem puzzling. We're supposed to secure lots of
 12 perceptual knowledge of apples. But I've argued that there is no appropriately distinctive
 13 way that visually identifiable apples actually *look*. This might tempt the thought that it can't
 14 really be perceptual knowledge we secure in these cases.

15 5.1. Limiting Visual Knowledge

16 One possibility is that in difficult cases like BOBBING APPLE things don't actually strike
 17 you as apples. Do they merely raise a suspicion; encourage investigation? But as a general
 18 rule this move is implausible. It's true of course that there will be many tricky situations
 19 where F 's don't strike you as such. But picture yourself at dusk in an overgrown field. You
 20 come across a dog-chewed, mud-stained little greenish thing half hidden in the grass. Surely
 21 we can imagine it striking you quite forcefully and spontaneously that the thing's a tennis
 22 ball. Taking this for granted you wonder if it's the one you lost.

23 Another possibility is that even if *Spin* does really strike you as an apple when you look –
 24 even if you get that 'just see' phenomenology (McGrath 2017: 1) – it's not perceptual
 25 knowledge you'll get to secure. If your knowledge here is not perceptual Looks Views
 26 needn't account for it. And after all, *Spin*'s context and your background experiences must
 27 be playing some ineliminable role in securing your knowledge in BOBBING APPLE. Does
 28 this make your knowledge here distinctively inferential or indirect?

1 One way or the other this is at least not obvious. There need be no difference in how it
 2 seems to spot that *Real* and *Spin* are apples, for example. Aspects of context sometimes are
 3 but often are not identifiable elements in how we might form either kind of judgement. At
 4 the same time is context *not* somehow at work even in REAL APPLE? Swap in only slightly
 5 more ascerbic lighting and perhaps *Real* could have struck you as a fake.

6 And while they're motivated by a puzzle there's something puzzling about these moves to
 7 restrict perceptual knowledge in return. For they imply that we *only* secure perceptual
 8 knowledge in the most benign and canonical situations – when the sun shines and we've
 9 time on our hands. But it would be remarkable to discover that that was how perception
 10 worked.

11 There are plenty of visually problematic, fleeting situations in which it would at least be
 12 *good* to secure some perceptual knowledge. Spotting the bad apple might lead you to the
 13 good ones; might confirm your suspicion that the herd had moved on a while ago. It surely
 14 pays to react quickly and decisively to the stalking predator without jumping at every
 15 shadow. Isn't this kind of *difficult* work – in the wild, in real time – something our perceptual
 16 capacities are honed to perform?

17 So the attempt to retrench here is best seen as a reaction to a sense of the *impossible*. If *not*
 18 via their look, then *how*?

19 5.2. Analogy with Perceptual Constancy

20 It's useful at this point to reflect again on the analogy we drew with our ability to see and
 21 track things' basic looks – their colours and shapes (§2.2). You can know what colour
 22 something is just by seeing its colour. And it's uncontroversial that your ability to get
 23 hooked up to things' colours in this way is both robust and fine-grained.

24 A red ball can look all sorts of ways and still look *just that shade of red*. You can track that
 25 shade of red as the light moves over the ball's polished glaze; can track it through jagged
 26 shadow and well into twilight. So your capacity to spot colours is robust. And you can
 27 distinguish differently coloured objects even when they stimulate your retina with the same
 28 frequencies of light – the bright red ball in shadow; the dull red apple in bright sun. The
 29 penny that's gone through the mangle at your local arcade looks elliptical; the legal tender
 30 held elliptically does not. The visual distinctions here are fine-grained. Where basic looks

1 are concerned such robust and fine-grained abilities manifest the familiar phenomena of
2 perceptual constancy.²⁹

3 That analogy delivers three lessons. Firstly it shows how you can get hooked up to some
4 stable features of your environment – things’ colours and shapes – in the absence of any
5 stable signal. Whatever the mechanisms that underpin this, it’s clearly not impossible. This
6 allows us to take seriously the idea that you might get hooked up to something’s *being an*
7 *apple* but not by acquaintance with any stable *look*.

8 Secondly it’s clear that the mechanisms which hook us up to things’ colours and shapes are
9 massively sensitive to context. Most visual illusions exploit the other side of this coin.
10 Bookended by differently oriented chevrons the two equal lines of the Muller-Lyer illusion
11 look different in length. With a brief enough glimpse the black and white line drawing of
12 the strawberry strikes you as slightly red (Macpherson 2012).

13 Given the paucity of information contained in the light array itself it’s unsurprising that
14 our perceptual mechanisms make use of all the context and experience available to them.
15 They have to filter out huge signal noise. And with the ephemeral signals as inputs, the
16 problem of identifying stable environmental features is after all *ill-posed*. What those rare
17 visual illusions manifest is how in spite of the cognitive hurdles our perceptual systems get
18 to be so reliable (cf. Lyons 2011).

19 So: visual knowledge even of things’ superficial visual appearances is sensitive to the details.
20 Being sensitive to context doesn’t set our visual knowledge of things like apples apart. In
21 fact it manifests something characteristic of our most uncontroversially perceptual abilities.

22 The third lesson we should draw from the analogy is that the complex cognitive machinery
23 that delivers our visual knowledge doesn’t automatically do duty as the normative grounds
24 of our visual knowledge.

25 In tracking something’s colour we can be aware both of the stable appearance and the
26 signal. But it would be odd to treat those thin and fleeting visual ephemera as anything like

²⁹ See e.g. Allen 2018; Cohen 2015.

1 our reasons for judging as we do; as our *evidence*.³⁰ The light array – causal intermediary
 2 though it is – is not stable enough to provide you with such reasons. Instead the signal here
 3 is something more like noise than evidence – something to be cut through as opposed to
 4 the basis upon which you know.

5 Something similar is going on when it comes to our visual knowledge of many rich features.
 6 Wherever we find that our ability to spot *F*s is robust and fine-grained in relation to the
 7 causal intermediaries we'll find that those messengers, whether basic looks or retinal
 8 stimulations, will fail as normative grounds for the knowledge we secure with their help.
 9 In this respect compare *artificial* perceptual systems. Our most robust and fine-grained
 10 models are deep learning systems. We have complete access to all their inputs and internal
 11 processes – the causal intermediaries between the target features and the system's
 12 identification of them. Nonetheless explanatorily they remain *black boxes*.³¹ Knowledge of
 13 the causal chain doesn't reveal why these systems are reliable; what accounts for the
 14 positive epistemic status of their outputs; how they deliver us knowledge.

15 In sum visual acquaintance with things' features is possible even when there's no stable
 16 intermediary. Our visual acquaintance with even sensory qualities is hugely sensitive to
 17 context. Finally, having a robust and fine-grained ability to visually discriminate *F*s makes
 18 it plausible that we can be visually acquainted with *F*s – that the causal intermediaries don't
 19 account for how we know about *F*s.

20 6. Intuitions about Looks

21 I've presented and defended an argument against Looks Views. But my conclusion may
 22 still have about it an air of the magical. I've said that we often cannot appeal to how things
 23 look in accounting for our ability to identify those things by looking. This can sound
 24 absurd. It's as if we're denying those common sense thoughts we began with (§2.1).

25 I'll take each in turn.

³⁰ It was a founding assumption of a long-since-rejected early modern empiricism that such ephemeral looks – sense data – would ground knowledge even of things' superficial visual appearances. See e.g. Locke 1690, 2.9.1. and 4.4.3. Very plausibly this restrictive model of our basic evidence is a root cause of that project's failure.

³¹ See e.g. Watson 2022; also Sullivan 2022 and Zednik 2021 for discussion.

1 6.1. Ways of Talking about Looks

2 It's plausible that if an object looks no way at all to you, you cannot secure perceptual
3 knowledge about it by looking.³² And any object that looks some way to you will interact
4 with light. So you can only visually identify *F*s if they have some basic looks.

5 Even so I've suggested (§4.2) that visually identifying apples needn't involve *acquaintance*
6 with their determinate, basic look. *Spin* looks some way to you – it stands out against its
7 backdrop. But you can't make out much of the pattern of *Spin's* colours; nothing of its
8 freckles.

9 It would also sound odd to say that you can tell by looking that something is an apple
10 without knowing what apples look like. But we can understand the claim that you know
11 what apples look like as the claim that you're good at spotting apples – that you really do
12 have this robust and fine-grained visual ability. What we've uncovered is that there need
13 be no perceptually privileged look which mediates that ability.

14 Another way to understand the claim that you know what apples look like is as the claim
15 that you've learnt what apples look like. It's plausible that as a process this might involve
16 knowledge of that classic apple look, *L_{CA}*. And early in the process perhaps your capacity
17 to recognize apples by looking didn't outrun your awareness of such canonical looks. But
18 to the extent that your ability to spot apples has become robust and fine-grained you have
19 outgrown the picture books.

20 6.2. Looks as Reasons

21 When you said that the thing looked like an apple it was tempting to treat you as offering
22 me *your reason* for thinking it an apple.³³ I've argued that looks can often not play this
23 evidential role. Luckily we can make sense of your words without interpreting you as trying
24 to present me with your evidence.

25 Depending on the timbre of the question to which 'it *looks* like an apple' is supposed to
26 provide an answer one possibility is that it is meant not as a justification but as a *retreat*.³⁴

³² While I don't need to doubt it here this is questionable. See footnote 12 above.

³³ See McGrath 2018, p117.

³⁴ Compare Gauker 2018, p158.

1 You might want to step back from claiming to know that the thing is an apple, especially
2 if I've raised my eyebrows.

3 In a less defensive spirit you might simply be telling me *how* you know. Discovering that
4 your belief is visual helps me to evaluate your testimony. I needn't worry about whether
5 smell or taste or others' words are unreliable round here. I can concentrate on whether
6 vision is safe enough or whether you really know what apples look like.

7 Framing things this way makes clear how you might provide others with reasons for belief
8 without giving them *your* reasons. In one way or another I take it that testimony is very
9 often like this. At the same time though you might intend by your retort something less
10 testimonial and more *demonstrative*. 'Just look at it!', you might signal: doing so should leave
11 us in no doubt.

12 6.3. Reasonable Mistakes

13 Finally, consider the case where you mistake a perfectly forged cake for an apple. Your
14 mistake here seems entirely reasonable. The locus of the problem is where the cake is – its
15 superficial visual appearance. Doesn't this make it clear that *looks* play some ineliminable
16 role in accounting for why, when they are, our visual judgements are reasonable?

17 The cake manifests a determinate pattern of basic looks. This look is a significant aspect
18 of the causal chain that results in its striking you visually as an apple. That specific
19 determinate look – like L_D – can't warrant your belief that it's an apple. You've never seen
20 it before; can't have associated it with being an apple. Yet any characteristic or common
21 look we might collar – anything like L_{CA} – won't be appropriately related to *being an apple*.

22 What then can we say about such cases? Reflect again on the contrast between this and the
23 hallucinogen-induced belief. Where the master forger's cake strikes you as an apple your
24 belief – like the beliefs of so many others – manifests a robust and fine-grained ability to
25 spot apples.³⁵ Securing empirical knowledge is an inherently risky enterprise. Occasionally
26 we can be tricked. But where hallucinogens are involved something about you is broken.

³⁵ This response bears some structural similarities to Schellenberg's *capacity first* view of perception (2018 especially chp2). For Schellenberg capacities are individuated in terms of their proper function. While I'm sympathetic to the project the thought here is that it's the robust and fine-grained *reliability* of our perceptual systems against which we can judge whether or not a perceptual belief is reasonable. One benefit of the current view is that it thereby gives us a tool to determine *which* features we have the visual capacity to spot.

1 In the circumstances you no longer manifest that robust and fine-grained ability for
2 spotting apples.

3 7. Conclusions

4 I've argued on normative grounds that no look will account for the perceptual knowledge
5 of *Fs* we secure when our capacity to visually identify *Fs* is robust and fine-grained. The
6 look would need to be shared by *Fs*' doppelgangers. And identifying the doppelgangers
7 involves working out which things, in which situations, actually strike us visually as *Fs*. Yet
8 often the pattern of dispositions we actually manifest is going to be too complex to be
9 accounted for by acquaintance with any superficial appearance.

10 Put things another way. Suppose we try to identify *what it is* about how you came to know
11 that this thing was an apple that accounts for how you know that. We either find things
12 idiosyncratic enough to this encounter that they do the causal but not the epistemic work
13 – they fail by ST2. Or we find things which are generic and stable but don't do the
14 normatively relevant causal work – they fail by ST1.

15 Nothing I've said implies that looks are *no* help. There are common, characteristic and
16 distinctive looks. We can think and talk about these, share in or argue over them. Their
17 stability certainly distinguishes looks from the light signal. Things' appearances provide us
18 with some important scaffolding. It's just that these identifiable looks don't often account
19 for why, in the real world, things actually strike us in the many rich ways they do. Given
20 this they can't account for how we thereby secure so much of our visual knowledge.

21 If we can sometimes secure visual knowledge of things' rich features but not *on the basis of*
22 acquaintance with their superficial visual appearances then we sometimes secure that
23 knowledge directly. We should conclude that we can be visually acquainted not only with
24 how things look but with at least some of those rich features we were looking for.
25 Appearances are not a veil.

26 Yet if the floodgates are open does this mean we can see not only things' being apples but
27 also for example the book's unhelpfulness or that angry person's need of a good lawyer?
28 We'd have reason to acquiesce only if our abilities to spot such features by looking were
29 robust and fine-grained. I take it that they are not.

BIBLIOGRAPHY

- Allen, K. (2009). 'Inter-species variation in colour perception'. *Philosophical Studies* 142(2): 197-220.
- _____ (2018). 'Perceptual constancy and apparent properties'. Dorsch & Macpherson (eds.). *Phenomenal Presence*. OUP. pp. 39-57
- Austin, J. L. (1961). "Other Minds". In J. L. Austin (ed.), *Philosophical Papers*. Oxford, England: Clarendon Press.
- Austin, J. L., & Warnock, G. J. (1962). *Sense and Sensibilia*. Oxford: Clarendon Press.
- Brewer, B. (2011). *Perception and Its Objects*. Oxford: OUP.
- Cassam, Q. (2007). 'Ways of Knowing'. *Proceedings of the Aristotelian Society* 107 (1pt3):339 - 358.
- Cohen, J. (2015). Perceptual constancy. *The Oxford handbook of philosophy of perception*, 621-639.
- Dretske (1969). *Seeing and Knowing*. Chicago: UCP.
- French, C. (2014). 'Knowledge and Ways of Knowing'. *Proceedings of the Aristotelian Society* 114 (3pt3):353-364
- Gauker, C. (2018). 'Do Perceptions Justify Beliefs? The Argument from "Looks" Talk'. In Gersel *et al* (eds.), *In Light of Experience: Essays on Reason and Perception*. Oxford: OUP. pp. 141-160
- Johnston, M. (2006). 'The Function of Sensory Awareness.' In Tamar Szabo Gendler & John Hawthorne (eds.), *Perceptual experience*. New York: OUP.
- _____ (2018) 'Sensory Disclosure: Neither a Propositional, Nor a Factive, Attitude'. In Grzankowski, & Montague (eds), *Non-Propositional Intentionality*. Oxford: OUP.
- Locke, J. (1690). *An Essay Concerning Human Understanding*. Ed. P. Nidditch 1975. Oxford: Clarendon.

- Lyons, J. (2011). 'Circularity, reliability, and the cognitive penetrability of perception'. *Philosophical Issues* 21 (1):289-311.
- Macpherson, F. (2012). 'Cognitive penetration of colour experience: Rethinking the issue in light of an indirect mechanism'. *Philosophy and Phenomenological Research*, 84(1), 24–62.
- Martin, M. G. F. (2010). 'What's in a look?' In B. Nanay (ed) *Perceiving the World*. Oxford: OUP. pp160-225.
- McGrath, M. (2017). 'Knowing what things look like.' *Philosophical Review* 126(1): 1-41.
- _____ (2018). 'Looks and Perceptual Justification.' *Philosophy and Phenomenological Research*: 110-133.
- McNeill, W. E. S. (2012). 'On Seeing That Someone is Angry'. *European Journal of Philosophy* 20(4): 575-597.
- _____ (2015). 'Inferentialism and our knowledge of others' minds. *Philosophical Studies*, 172(6), 1435-1454.
- Millar, A. (2000). 'The Scope of Perceptual Knowledge.' *Philosophy* 75(291): 73-88.
- _____ (2011). 'How visual perception yields reasons for belief.' *Philosophical Issues* 21(1): 332-351.
- Peacocke, C. (1983). *Sense and Content: experience, thought, and their relations*. Oxford, Clarendon Press.
- Schellenberg, S. (2008). 'The situation-dependency of perception.' *The Journal of Philosophy*, 105(2): 55-84.
- Schellenberg, S. (2018). *The unity of perception: Content, consciousness, evidence*. Oxford University Press.
- Shieber, J. (2017). 'Looks and the Immediacy of Visual Objectual Knowledge.' *Analysis* 77(4): 741-750.
- Smith, J. (2015). 'The Phenomenology of Face-to-Face Mindreading.' *Philosophy and Phenomenological Research* 90(2): 274-293.

- Sullivan, E. (2022). ‘Understanding from Machine Learning Models.’ *British Journal for the Philosophy of Science* 73(1):109-133.
- Tucker, C. ed. (2013). *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*. OUP: Oxford.
- Travis, C. (2004). ‘The Silence of the Senses.’ *Mind* 113 (449):57-94.
- Watson, D. S. (2022). ‘Conceptual challenges for interpretable machine learning’. *Synthese*, 200(1), 1-33.
- Zednik, C. (2021). ‘Solving the black box problem: A normative framework for explainable artificial intelligence.’ *Philosophy & Technology*, 34(2), 265-288.