

Looks Non-transitive!

PHILIPPE CHUARD

Philosophy Program, Research School of Social Sciences,
Australian National University
philippe@coombs.anu.edu.au

RICHARD CORRY

School of Philosophy,
University of Tasmania, Hobart
richard.corry@utas.edu.au

Suppose you are presented with three red objects. You are then asked to take a careful look at each possible pair of objects, and to decide whether or not their members look chromatically the same. You carry out the instructions thoroughly, and the following propositions sum up the results of your empirical investigation:

- i. red object #1 looks the same in colour as red object #2.
- ii. red object #2 looks the same in colour as red object #3.
- iii. red object #1 does not look the same in colour as red object #3.

Experiences like these may seem familiar. Had you recently been in a hardware shop to choose a suitable shade of beige for repainting your kitchen wall, you might have experienced something similar. The shades of beige on two colour chips are indiscriminable. Yet, when you compare them with a third colour chip, while the latter is indiscriminable from the shade on the second chip, it is discriminable from the first.¹

Such experiences suggest that the relation denoted by the predicate *looks the same as* isn't transitive. That is, the following thesis is not true of *looks the same as*:

(TRAN) $(\forall x) (\forall y) (\forall z) [(x \text{ looks the same as } y) \ \& \ (y \text{ looks the same as } z)] \rightarrow (x \text{ looks the same as } z)$.

¹Typically, talk of “indiscriminability” is supposed to be equivalent to talk of “looking the same as”—for the kind of indiscriminability at play here is purely perceptual. In the rest of the paper, however, we follow Delia Graff's (2001: 911-12) focus on the latter expression. For Graff (*ibid.*), “looks the same as” is to be understood in its phenomenal and non-dispositional sense. The relation denoted holds between (actually) perceived objects—not

To be sure, all it takes to falsify (TRAN) is a situation that instantiates:

(INT) $(\exists x) (\exists y) (\exists z) (x \text{ looks the same as } y) \ \& \ (y \text{ looks the same as } z) \ \& \ \neg (x \text{ looks the same as } z)$.

And the sort of experience(s) described above seem to provide compelling reasons to believe that *looks the same as* is not transitive. Of course, there may be other experiences such that all the objects perceived are indiscriminable from one another. Still, a few experiences suffice to make (INT) true.

This line of thought has very much been orthodoxy, at least since Nelson Goodman's (1951) *The Structure of Appearance*. Yet, a small number of philosophers have resisted it—Frank Jackson and R.J. Pinkerton (1973) appear to have been the first to dissent.² Recently, Delia Graff (2001) has critically reviewed the motivations advanced against (TRAN). She concludes that (TRAN) has to be true on the ground that “while there are very good reasons to think it is, there is not sufficient reason to think it is not” (2001: 932).³ About the experiences alleged to support (INT), she writes:

I want to incite a bit of doubt that we really do have such experiences, but I also argue that our having such experiences (if we do have them) does not provide the required support for non-transitivity, [...]. (Graff, 2001: 906)

between perceptual experiences of these objects—, and must be relativised to particular perceivers in particular circumstances. For a detailed comparison of different relations of indiscriminability, see Williamson (1990).

² See also Hardin (1988), Jackson (1977: 113-15), Mills (2002), and Raffman (2000). Williamson (1990: 12-14) contains a critical discussion of some of the arguments advanced—see also Graff (2001: 914-5).

³ At first sight, this is somewhat surprising, since the claim that *looks the same as* is non-transitive might seem easier to justify than the claim that it *is* transitive. After all, all it takes to justify (INT) is just one instance of three experiences (or more) such that:

$\neg (\forall x) (\forall y) (\forall z) [(x \text{ looks the same as } y) \ \& \ (y \text{ looks the same as } z)] \rightarrow (x \text{ looks the same as } z)$.

The justification of (TRAN), on the other hand, would seem to require that one establishes that, for *any* experience:

$\neg (\exists x) (\exists y) (\exists z) (x \text{ looks the same as } y) \ \& \ (y \text{ looks the same as } z) \ \& \ \neg (x \text{ looks the same as } z)$.

How the latter can ever be established, however, is far less straightforward.

Graff outlines two kinds of consideration in favour of (INT) and argues that, in fact, each fails to support it.⁴ The first exploits the fact that our discriminatory powers are limited: as a result, a succession of indiscriminable variations in a certain respect might appear to make (INT) true. Call this the ‘argument from discriminatory limitations’. The second consideration concerns the existence of phenomenal continua, and whether parts of such continua look homogeneous—the ‘argument from phenomenal continua’. Graff raises a variety of problems for these two considerations and concludes that (INT) is ill-motivated. In this respect, she attempts to shift the burden of proof onto the defenders of (INT): despite what has been assumed for the past fifty years, there is no good argument in favour of the orthodox view—or so she argues.

We remain unconvinced. Of course, (INT) is often taken for granted these days. In this regard, Graff’s attempt to shift the burden of proof doesn’t seem like a bad idea. But we doubt whether she succeeds, as we find Graff’s objections far from cogent. Some of these objections underline what we take to be fairly superficial problems, to be fixed without much ado. Others appear more substantial in that she advances radical re-descriptions of the sort of experiences alleged to support (INT). But we find such re-descriptions incorrect—at best inconclusive. The thrust of this paper is thus entirely negative. Familiar considerations offered in support of (INT), we argue, are left unscathed by Graff’s objections. We don’t defend (INT) in any straightforward manner. But if you find the considerations advanced in favour of (INT) compelling, we try to show why Graff’s objections should not make you change your mind.

The paper has three parts. Section 2 discusses Graff’s main objection to the ‘argument from discriminatory limitations’, while section 3 reviews some of the difficulties which, according to Graff, plague the ‘argument from phenomenal continua’. Before this, we take a brief look at Graff’s motivation for attempting to resurrect the claim that *looks the same as* is transitive, and we suggest that, in fact, she offers no good reason to stick with (TRAN).

⁴ A third consideration—‘Wright’s proof’, from Wright (1975)—is described by Graff as a more formal version of the second kind of consideration, and we shall ignore it here.

1. Why not (INT)?

Why go against orthodoxy and argue that *looks the same as* is transitive? Graff's reason for rejecting (INT) is the need to defend what she calls 'Same Appearance Claims':

(SAC) if x and y look the same with respect to colour, then (if x looks red, y looks red too).

She takes such claims to be truisms: surely, she reasons, if two things look the same with respect to their colour, it ought to be the case that if one thing looks to have a particular colour, the other should look to have that very same colour too. Otherwise, she insists, they wouldn't look chromatically the same in the first place (Graff, 2001: 909)!⁵

But why does a truism like (SAC) need defending? Because, says Graff (2001: 907), it plays a crucial role in phenomenal sorites arguments, and such arguments are supposed to raise a paradox—although valid, they appear to have true premises and a false conclusion. In brief, the threat phenomenal sorites arguments place upon (SAC) owes to some of the traditional ways in which such a paradox has been addressed. For instance, in order to resolve the paradox, many have been drawn to deny the central premise in such arguments: (SAC). Since she takes such a claim to be a truism, Graff is obliged to offer an alternative diagnosis of what goes wrong with phenomenal sorites arguments. She suggests that, contrary to what is normally assumed, the sort of examples exploited by sorites arguments does not instantiate the non-transitivity of *looks the same as*. However, granted that Graff's suggestion offers a plausible solution to the paradox, it does not entail that *nothing* instantiates (INT)—or so we argue.

1.1. *Phenomenal Sorites*

But first, here is how the problem arises in a little more detail. Two main ingredients are needed to construct a phenomenal sorites argument—or, indeed,

⁵ Instances of (SAC), as Graff notes (2001: 907), need not be restricted to visual experiences: (SAC) may cover "appearance" claims in other sensory modalities too.

any sorites argument.⁶ The first is a *sorites series*: a series of individuals, ordered with regard to their close similarity—but not complete identity—in a given respect. Such a series might contain middle-aged men, ordered in terms of their height; or perhaps, the same men, but in a different order, relative to the number and distribution of hair on their skull; or collections of grains of sand, quantitatively ordered. In the phenomenal case, a classic example is a series of coloured patches, ordered in terms of the continuous variations in the colour they *appear* to have.

Such series, in turn, need to instantiate two related features. First, *indiscriminability*: each member of the series needs to be closely similar to adjacent members in the relevant respect—in fact, so similar that a normal subject couldn't discriminate between them in that respect. For instance, the distribution and number of hairs on the skull of middle-aged men might vary by just one hair from a member of the series to the other. Of course, these two men differ in that one has one more hair than the other. The difference, however, is presumed to be insignificant. In the case of phenomenal sorites series, adjacent members have to *look the same*—not just similar—in the relevant respect, as when two coloured patches with a similar, but slightly distinct, shade of red look chromatically the same.

Second, *non-transitivity*: as we shall see, it is important that phenomenal sorites series be such that the relation of *looking the same as*, which holds between each adjacent member on the series, is non-transitive. It is not sufficient that every patch in the series looks the same as its closest neighbours. A series of patches of *exactly the same colour* could instantiate that. What is required is that the patches in question vary in a minute way—so minute that the variation is indiscriminable—with respect to their colour, to the effect that the patches at each end of the series look radically different.⁷

⁶ For more on these arguments, see e.g. Burns (1991: ch. 4 & 5), Read (1994: ch. 7), Sainsbury & Williamson (1997), Williamson (1994). For similar descriptions of the conditions required to construct a sorites argument, see Barnes (1982), Hyde (2002).

⁷ At least, this is the traditional way of presenting such series. But note that the assumption that the patches actually differ chromatically is perhaps unnecessary. Consider, for instance, Chevreul's array (see Clark, 1993: 14), where a strip of uniform grey is presented with a lighter strip of grey on its left, and a darker one on its right. Every part of the central strip of grey has the same colour. Yet, due to the different strips of grey alongside the

The second element required to construct a sorites argument is a sorites premise. For instance, in the case of a sorites argument with bald middle-aged men, the sorites premise is:

(BALD) $(\forall x)(\forall y)$ if x is a middle-aged man with n number of hair (distributed in such-and-such way) and y is a middle-aged man with $n + 1$ hair), then (if x is bald, y is bald too).

With phenomenal sorites arguments, (SAC) provides the sorites premise. Note the presence of a *vague* predicate in such premises: while the predicate clearly applies to one part of the series and clearly fails to apply to another part, it admits of borderlines cases. That is, there is a point on the series at which it is unclear whether or not the predicate applies. Perhaps, there is no fact of the matter whether it applies. Or perhaps there is, but we can have no knowledge of that fact.⁸ Typical examples include ‘is bald’, ‘is tall’, ‘makes a heap’, and so on. With phenomenal sorites series, ‘looks red’ is one such predicate. As Graff (2001: 907) observes, such a predicate must be understood in its phenomenal sense: ‘looks red’ constitutes a purely *observational* predicate, so that “its applicability to an object (given a fixed context of evaluation) depends only on the way that object appears” (*ibid.*). Thus, ‘looks red’ clearly applies to the patches at one end of the series, since they *are* red; and it clearly doesn’t apply to patches at the other end, which aren’t red. But it is unclear at which point in the series (if any) the patches stop being red.

The argument then takes the following form—assume we have a phenomenal sorites series of thirty patches, ranging from red to yellow:⁹

central strip, the left-hand side of the central strip looks darker than its right-hand side. We could thus divide the central strip of grey into very many sub-portions, so as to construct a phenomenal sorites series. Each adjacent small sub-sections are indiscriminable, but the extremities of the central strip of grey look different.

⁸ How to best characterise vague predicates and their so-called ‘borderline cases’ is a matter of some controversy: see Greenough (2003). For a general discussion of different accounts of vagueness, see Burns (1991), Merricks (2001), and Williamson (1994).

⁹ For the various forms sorites arguments can take, see Burns (1991: ch. 4), Sainsbury & Williamson (1997), and Hyde (2002).

- (1) patch #1 looks red.
- (2) $(\forall x)$ if (patch # x looks the same as patch # $x+1$), then (if patch # x looks red, patch # $x+1$ looks red).

Therefore,

- (3) patch #30 looks red.

The interest surrounding sorites arguments owes to the paradox they present. Once the auxiliary premises are made explicit, such arguments are clearly valid. Furthermore, all their premises seem to be true. Yet, the conclusion is obviously false. In the example above, patch #30 *is* in fact yellow, and thus should look yellow—*contra* what (3) says. However, according to the classical notion of validity, there cannot be valid arguments with true premises and a false conclusion.

In order to solve such paradox, the main reaction—at least with non-phenomenal instances of sorites arguments—is to reject one of the premises. The sorites premise, in particular, is the prime target of such a response. Graff concurs (2001: 908): she thinks that a one-hair difference between two middle-aged men is in fact sufficient to make it the case that, while one is bald, the other isn't.

But what about phenomenal sorites arguments? Should the sorites premise be rejected too? We are now in a position to understand why Graff thinks that (SAC) needs defending. Since (SAC) *is* the sorites premise for phenomenal sorites arguments, and since Graff finds (SAC) to be a non-negotiable truism, the denial such a premise is simply not an option for her. She must find something else. Among the various features needed to construct a phenomenal sorites argument, she picks on the phenomenal sorites series and, in particular, the fact that it instantiates (INT):

It is not so obvious, however, that there could be (and hence that it is legitimate to suppose that there could be) a sorites series of colour patches on which *looking the same as* is not transitive. That such a supposition forces the retraction of same-appearance claims, which really do seem like truisms, provides grounds for calling the supposition into question. (Graff, 2001: 909)

We think that this is a bit of an overreaction on her part. For one thing, contrary to what Graff appears to suggest, it doesn't look as though the non-transitivity of *looks the same as*, in and of itself, "forces the retraction" of (SAC). The latter, it seems

to us, is compatible with non-transitivity. For another, even if we accept Graff's proposed solution to the phenomenal sorites, such a solution does not warrant the general rejection of (INT). Graff fails to distinguish between the different things that might instantiate (INT). Thus, while she denies that phenomenal sorites series instantiate the non-transitivity of *looks the same as*, she does not appreciate that such a rejection is compatible with the fact that (INT) applies to other things—like parts of sorites series, for example. We take each of these points in turn.

1.2. *Non-transitivity & Same Appearance Claims*

Graff writes as if, once one endorses (SAC), one is thereby committed to the rejection of (INT), because the combination of these two claims entails a contradiction. It's unclear why she might think such a thing, since, at first sight, (INT) seems compatible with (SAC).¹⁰ Focus on just three patches of a phenomenal sorites series and suppose they instantiate the non-transitivity of *looks the same as*. Thus:

- i) patch #5 looks chromatically the same as patch #6.
- ii) patch #6 looks chromatically the same as patch #7.
- iii) patch #5 does *not* look chromatically the same as patch #7.

The situation described by (i), (ii) and (iii) suffices to make (INT) true. And it seems compatible with (SAC), insofar as none of the following instances of (SAC) are false:

- iv) if patch #5 looks chromatically the same as patch #6, then (if patch #5 looks red, patch #6 looks red too).
- v) if patch #6 looks chromatically the same as patch #7, then (if patch #6 looks red, patch #7 looks red too).

¹⁰ A *caveat*: This point holds, insofar as, following Graff, (SAC) is phrased in terms of observational predicates such as “looks red” or “looks yellow”—and even with more precise predicates like “looks red₁₃” or “looks red₅”. The point may not hold, however, when a more precise colour predicate such as “looks exactly the same colour as patch #1 looks” is substituted in (SAC). But even then, one cannot straightforwardly derive a contradiction from the combination of (SAC) and (INT), without resorting to some additional assumptions. Even so, there are at least two ways to resist such a conclusion. We discuss this point elsewhere, as we lack the space to do so here.

vi) if patch #5 looks chromatically the same as patch #7, then (if patch #5 looks red, patch #7 looks red too).

First of all, we take it that both (iv) and (v) are uncontroversially true: both conditionals have true antecedents and true consequents. Second, the conditional in (vi) is true too: it has a false antecedent—patch #5 doesn't look the same as patch #7. And so the whole conditional is true, regardless of the truth-value of its consequent—that is, assuming a classical conception of the semantics of conditionals. But it's worth remarking that its consequent is true too—even though they look chromatically different, it is still the case that both patch #5 and patch #7 look red. In fact, even if one looked scarlet and the other purple, they would still both look red. The observational predicate 'looks red' is determinable and thus covers a wide range of different shades of red.

This shows, it seems to us, that the non-transitivity of *looks the same as* need not force the retraction of 'Same-Appearance' claims *per se*—a point that Graff fails to appreciate.

1.3. *Phenomenal Sorites Series & their Parts*

Furthermore, it seems to us that Graff's alternative solution to the paradox presented by phenomenal sorites arguments offers no reason to reject (INT) in general. Graff's solution seems to amount to this claim: deny that there are phenomenal sorites series, the *whole of which* instantiates the non-transitivity of *looks the same as*. That is, she denies that there exists a series of thirty colour patches such that, while (i) every two adjacent patches on the series are chromatically indiscriminable, (ii) the end points of the series look to have different colours. In other words, there must be at least one cut-off at some point on the series, to the effect that two adjacent patches look chromatically different.¹¹ The existence of such a cut-off suffices to make one of the conditional premises of a phenomenal sorites false. In which case, the paradoxical situation dissolves, since the argument is unsound after all.

But does this really amount to an outright denial of the non-transitivity of *looks the same as* in general? Of course not. (INT) says that *there are* experiences such that,

¹¹ See her discussion at p. 913.

when object #1 looks the same as object #2 in some respect ϕ , and object #2 looks the same as object #3 in the same respect ϕ , it is not the case that object #1 looks the same as object #3 ϕ -wise. To deny that phenomenal sorites series instantiate such a pattern is *not* to say that *no* collection of experienced objects ever instantiate it either. In fact, it might even be the case that, although the whole phenomenal sorites series doesn't instantiate (INT), parts of it do—such as patch #5, patch #6, and patch #7, for instance. Therefore, insofar as Graff only needs the claim that whole phenomenal sorites series do not instantiate the non-transitivity of *looks the same as*, she provides no good reason to deny that *looks the same as* is non-transitive.

On this ground, we don't think that Graff successfully motivates her rejection of (INT). In what follows, we argue that she does not provide any good reason either for thinking that the traditional considerations in favour of (INT) are ineffectual.

2. Discriminatory Limitations

Now that we've seen why Graff wishes to reject (INT), we can take a closer look at her objections against the usual considerations advanced in favour of such a thesis. The first argument Graff considers rests on the idea that our capacity for perceptual discrimination is limited, such that “if a thing changes in an observable quality (e.g., *shape*) just by some sufficiently small amount, we will be unable to perceptually detect the change” (2001: 916). Graff suggests that the limitations in question can be cashed out in two different ways (2001: 917):¹²

- (A) For some sufficiently slight amount of change (in colour, sound, position, etc.), when we perceive an object for the entirety of an interval during which it changes by less than that amount, we perceive it as not having changed at all.
- (B) For some sufficiently slight amount of change (in colour, sound, position, etc.), we cannot perceive an object as having changed by less than that amount unless we perceive it as not having changed at all.

¹² We follow Graff's wording of these two claims. But we suspect that a better formulation of (A) is available, so as to avoid the difficulties that arise later in the argument (see below).

The crucial difference between (A) and (B) is that, while (A) describes a relation between experience and the world—the way our experiences are responsive to certain changes in the world, (B) characterises the kind of experience we can have—the way things are represented in experience. Having outlined these two different kinds of discriminatory limitations, Graff attempts to show that neither gives us a reason to believe that *looks the same* is not a transitive relation.

Here is a brief outline of her strategy. First, Graff admits that, were (A) to be true, it would indeed follow that *looks the same* is not transitive. However, she argues that, as it stands, (A) is clearly false; and that any attempt to reformulate (A) so as to avoid obvious falsehood will result in the collapse of (A) into (B). As for (B), Graff points out that, by itself, (B) does not give us any positive reason to believe that *looks the same* is not transitive. In order to get from (B) to non-transitivity, she says, we need to add an additional premise—namely, that changes in the observational quality under consideration appear continuous. She then argues that we have no good reason to believe this extra premise.

We consider Graff's objections to the claim that some changes appear continuous in section 3. In the current section, we focus mainly on the first part of Graff's strategy. We argue that the idea behind (A) can be formulated in such a way that (i) it is not obviously false, (ii) it does not collapse into (B), and (iii) it still leads to the non-transitivity of *looks the same as*.

2.1. From (A) to Non-transitivity

To see how (A) could lead to (INT), suppose that (A) provides a correct description of Winifred's limited powers to visually detect changes in some observational quality \mathcal{Q} . That is, suppose that there is some amount of change ε such that, when the quality \mathcal{Q} of an object in Winifred's perceptual environment changes by less than ε , she perceives the object as not changing at all with respect to \mathcal{Q} .¹³

Now, imagine an object o that changes by $\frac{3}{4}\varepsilon$ between t_1 and t_2 , and changes by a further $\frac{3}{4}\varepsilon$ between t_2 and t_3 , so that, from t_1 to t_3 , the object changes by a total of $1\frac{1}{2}\varepsilon$. Assume that o remains in Winifred's visual field for a period including t_1 , t_2 ,

¹³ For the sake of brevity, we will cease to refer explicitly to quality \mathcal{Q} : from now on, any mention of changes in an object are meant to refer to changes in the quality \mathcal{Q} of that object; and mention of changes in the appearance of the object are meant to refer to changes in the appearance of the object with respect to \mathcal{Q} .

and t_3 —to avoid further complications, we may further assume that Winifred’s perceptual attention is focused upon o at all times, and that she is particularly attentive to what happens regarding \mathcal{Q} . Since o changes by less than ε from t_1 to t_2 , it will appear to Winifred as not having changed at all. This implies that o will look the same to Winifred at t_1 as it does at t_2 . Similarly, o will look the same to her at t_2 and t_3 . Between t_1 and t_3 , however, the object has changed by an amount greater than ε .

At this point, it seems that we can safely assume that the object is likely to look different at t_1 than it does at t_3 . For if Winifred is like you and me, she must be able—at least sometimes, in some circumstances—to detect sufficiently large changes in an object o with respect to a quality like \mathcal{Q} . In such cases, of course, ε must be such that a change of $1/2 \varepsilon$ is likely to be noticed by Winifred.¹⁴ And so, the following is true: (i) o at t_1 looks the same as o at t_2 , and (ii) o at t_2 looks the same as o at t_3 , but (iii) o at t_1 does not look the same as o at t_3 . And this is just the situation needed to establish that *looks the same as* is not transitive.

2.2. *Why (A) is false*

So what is the problem with (A)? According to Graff (2001: 918-19), (A) is implausible because it makes us infallible detectors of non-change. In Graff’s initial formulation, (A) expresses a relation between the world and our experiences, such that objects which change by less than ε are perceived as not-changing at all. The problem arises most distinctively in cases where an object does not change. Since it doesn’t change, it is trivially true that such an object changes by less than ε . (A) entails that *whenever* an object in our perceptual environment does not change, we will *always* perceive it as not changing.

The problem with this consequence is that the human perceptual system sometimes falls prey to perceptual illusions. Stationary objects can appear to

¹⁴ Strictly speaking, (A) implies (INT) if the following two additional premises are true:

(P1) if an object in a subject S ’s visual field changes with respect to quality \mathcal{Q} by an amount greater than ε , then the object will appear to S to have changed.

(P2) some objects change with respect to quality \mathcal{Q} by amounts less than ε for short periods of time, although their total change over longer periods of time add up to a change greater than ε .

While (P1) states that sufficiently large changes can be perceived, (P2) makes a claim about the actual occurrence of such changes.

quiver, sticks appear to change shape when half submerged in water, and so on. But (A) makes such occurrences impossible. For it is formulated in such a way that changes smaller than ε and non-changes are both perceived *as* non-changes. Thus, (A) incorrectly rules out the possibility of perceptual illusions in which, for example, we wrongly perceive non-changes *as* changes. In its present formulation, it requires that perception of non-change is always veridical.

In a way, one cannot help but feel that the difficulty Graff outlines here is rather superficial: fixing (A)'s formulation should get us out of trouble! Of course, it will not do to reformulate (A) so that only sufficiently small, but *non-zero*, amounts of change are perceived as no change. If it is possible to have an illusory experience of change when an object is not changing, it is also possible to fall prey to the same illusion even when the object is in fact changing—whether or not the change is smaller than ε . The real source of the problem is that (A) has the form: “whenever we are confronted with a stimulus of type *S* we will have an experience of type *e*” (Graff, 2001: 919). Such a claim rules out the intuitive suggestion that it is possible to have any kind of experience, no matter what scene is before us.

The challenge, then, is to find a reformulation of (A) that is not so strong as to rule out the possibility of perceptual illusions of change *by fiat*, yet leads to the non-transitivity of *looks the same as*. Such a challenge isn't so hard to meet, however. The key is to note that a defender of (INT) need not claim that transitivity fails in *every possible* situation: only that transitivity *sometimes* fails. After all, a single instance of transitivity failure is enough to make a relation non-transitive. Thus, there is no need for (A) to apply to all possible experiences. It will be enough if it applies to a particular range of experiences. And presumably, veridical experiences should do.

So, if the trouble arises with *illusory* experiences, the obvious suggestion is that (A) be reformulated so as *not* to apply to cases of illusion. Graff (2001: 920) herself offers another formulation of (A) along these lines:

(A*) For some small amount of change, whenever an object changes by any lesser amount, we cannot have a *veridical* experience of that change.

Does this mean that we answered the challenge raised by Graff's objection to (A)? Not quite.

2.3. *The Threat of Collapse*

Graff thinks that another type of difficulty plagues this new formulation of (A). Having put (A*) on the table, she immediately paraphrases it as follows (2001: 920):

(A**) For some small amount of change ε , we cannot *visually represent* a change of any positive amount δ less than ε .

She then observes, quite correctly, that (A**) is equivalent to the claim that our powers of discrimination are limited in the sense captured by (B). Like (B), (A**) is about the way things are *represented* in experience, not about the relation between our experiences and the world. It imposes a restriction on visual representation and the kinds of experiences we can have. And so, Graff concludes, the foregoing attempt to rescue (A) simply collapses into (B).

Of course, the collapse of (A) into (B) would not present any problem, if (B) did provide a good reason to think that (INT) is true. But, as Graff goes on to argue, (B) supports (INT) only when it is coupled with what she regards as an implausible assumption about the existence of phenomenal continua. In brief, the problem with (B) is that it is compatible with the possibility that all perceived changes are discrete, according to Graff (2001: 919). (B) requires that no change, no matter its scale, can be perceived—veridically or not—as being less than ε . However, a change that is perceived as being more than ε might be discrete, as when a growing object appears to be taking a sudden ‘jump’ in its height. Insofar as the change is perceptually represented—correctly or not—as *being more than ε* , it is compatible with (B).

So what is wrong with discrete jumps? Graff does not offer any detailed argument in answer to this question, but the thought seems to run as follows. Suppose that o looks different at times t_1 and t_3 . If changes in appearance always involve discrete jumps, then such a jump must have taken place between t_1 and t_3 . Consider any time t_2 between t_1 and t_3 . If the jump took place at or before t_2 , o will look different at t_1 than it looks at t_2 . If the jump occurs after t_2 , then o will look different at t_2 and t_3 . Either way, we have no instance of (INT) where o should (i)

look the same at t_1 and t_2 , as well as (ii) at t_2 and t_3 , yet look different at t_1 than it does at t_3 .¹⁵

On this ground, Graff concludes, (B) only leads to failures of transitivity under the additional assumption that there exist phenomenal continua—that is, non-discrete changes in perceptual experience. (B) alone doesn't entail (INT). Furthermore, Graff argues (2001: 922-3), the friends of (INT) haven't properly motivated this additional assumption about phenomenal continua.¹⁶ And so, provided that (A*) and (A**) are equivalent to (B), neither provides a good reason for (INT). In which case, (A) cannot be reformulated in a way that (i) avoids the difficulties alluded to earlier, and (ii) supports (INT).

2.4. *Why the Collapse?*

But is (A*) really equivalent to (B), as Graff claims? We have some doubts. In particular, it is unclear why Graff believes that (A**) is just (A*) “put another way” (2001: 920). For one thing, the notion of ‘veridical experience’ in (A*) isn't co-extensive with the notion of ‘perceptual representation’ which figures in (A**). Obviously, a perceptual representation can be false—as in cases of illusory experiences. Hence, the latter notion has a broader extension than that of ‘veridical experience’, since it covers non-veridical cases too.

For this reason, (A**) is also stronger than (A*). While the latter does indeed explicitly restrict the claim that our perceptual discriminatory capacities are limited to cases of veridical experience, it also allows that there could be various kinds of illusory experiences of changes smaller than ε —it's just that (A*) has nothing to say about such cases. (A**), on the other hand, suggests that perceptual representation of any change *as* smaller than ε isn't possible. Of course, this entails

¹⁵ Actually, the conclusion is a little hasty, since the argument overlooks the possibility of numerous jumps in appearance, back and forth between t_1 and t_2 , and again between t_2 and t_3 . One might think that such jumps could occur in such a way that they ‘cancel each other out’, as it were. If so, an instance of (INT) could still be obtained, even with discontinuous changes. However, we don't hold much hope for such a strategy. For the sake of argument, we will grant that a more complex version of argument could be produced to allow for any number of jumps—though we are less clear about what happens if an infinite number of jumps is allowed.

¹⁶ In section 3, we consider Graff's argument that we have no reason to believe that there are such phenomenal continua. If, as we argue, the argument fails, there is one motivation for (INT) left standing by Graff's strategy: even if (A**) is equivalent to (B), it does entail (INT) in cases involving continuous changes.

that we cannot have veridical perceptual representations of changes smaller than ε —the desired result. (A**) also allows for the possibility of perceptual illusions of small changes experienced as bigger than ε . But it rules out a certain kind of illusion, since it entails that no illusory representation of any change *as* being smaller than ε is possible. Admittedly, this is much weaker than the problematic consequence derived from (A). Nonetheless, it is sufficient to differentiate (A**) from (A*).

Finally, (A**) is very much like (B)—as Graff points out—in that both propositions make claims about the nature of the kind of experiences we can have. (A*), on the other hand, seems more like (A): it concerns the connection between experience and the world.

So why does Graff slide from (A*) to (A**)? Our best guess is that she believes that the inability to form a perceptual representation of a change smaller than ε is the best explanation for why one cannot *veridically* experience such a change. In other words, she seems to assume that (B) is the most plausible explanation for (A*). Indeed, Graff fleshes out (A*) as follows:

If, for example, a thing grows by just a nanometer, we cannot have a veridical experience of that growth, for an experience as of a nm growth is not the kind of experience we could have. (Graff, 2001: 920)

The first half of this sentence—up to and including the word ‘growth’—describes the kind of limitation rendered by (A*), in terms of a world-experience relation. The second half of the sentence looks like an attempt to explain such a limitation in terms of a restriction on the kind of experience we can have *à la* (A**)—or (B).

We agree that this kind of explanation is quite natural. But Graff is not warranted in the move from (A*) to (A**), unless (A**) *is* the *only* explanation—or at least the only *plausible* explanation—for (A*). In what follows, we will suggest that there is another possible explanation for (A*): an explanation which seems as plausible as the one favoured by Graff. Furthermore, the explanation advanced by Graff is one that at least some proponents of (A) are likely to resist. In this respect, Graff’s paraphrase of (A*) in terms of (A**) somewhat begs the question against those theorists who believe that there is only one kind of limitation on our discriminatory powers—the kind of limitation captured by (A) or (A*).

2.5. Veridical Experiences

When Graff speaks of veridical experiences, she seems to have in mind experiences that represent the world correctly. On this view, the fact that a perceptual experience truly represents an object or a scene is both necessary and sufficient for it to be veridical. If this is what it means for an experience to be ‘veridical’, then a perceiver’s inability to have *veridical* experiences of a certain kind of change in the world can only be explained by the fact that her experiences cannot *represent* such a change. Indeed, this is exactly what Graff’s move from (A*) to (A**) takes for granted.

Such a conception of what it takes for an experience to be veridical is congenial to an understanding of perceptual discriminatory limitations in the sense captured by (B). Like (B), it ascribes a central role to the fact that experiences are representational. But would proponents of (A) share such a conception of veridical experience? Not all of them, surely.

It is worth keeping in mind that, as Graff points out (2001: 918), (A) and (B) are logically independent. In particular, it is possible that (A) or (A*) be true without (B) also being true. For instance, it may be that, given (A*), we simply cannot veridically perceive changes smaller than ϵ , although we sometimes misperceive changes bigger than ϵ as being smaller than ϵ —*contra* (B). Of course, this scenario might seem rather odd. Indeed, it seems quite credible that if there are limitations upon our discriminatory powers, such limitations are correctly characterised by both (A*) and (B). Nevertheless, given the logical independence between (A*) and (B), it is possible for one to claim that *only* limitations of the (A)-kind constrain our perceptual discriminatory capacities. On such a view, (B) is false, but not (A*). We can grant, for the sake of argument, that this is at least a possible view, which occupies some portion of logical space among those who endorse (A) and (A*).

One reason for endorsing such a view might have to do with the denial that perceptual experiences are representational states. Note that, unlike (B), (A) and (A*) need not be concerned with perceptual representation. Rather, the existence of perceptual discriminatory limitations is solely a matter of the relationship between experiences and the world, on such a view. Hence, a proponent of (A*) might think that perceptual experiences aren’t representational at all: they have no

content.¹⁷ This would explain why (B) is false, on such a view: there are no limitations on what can be represented in experience, because nothing is represented in experience. Such a theorist would not be too impressed by Graff's eagerness to paraphrase (A*) in terms which presuppose something like (B)—and a conception of experiences as representational.

So, at this point, the question is whether any other explanation for the truth of (A*)—which doesn't presuppose (A**)—is available? For Graff's move from (A*) to (A**) is warranted only if the latter is indeed the only way to account for the former. Thus, we need to investigate whether it is possible to account for certain limitations upon what can be *veridically experienced*, without postulating limitations upon what can and cannot be *represented* in experience. We would like to suggest that there is indeed a way to cash out the notion of 'veridical experience', tailored to the view that the only limitations on our discriminatory powers are those described by (A*). In which case, there is a clear way to distinguish (A*) from (B).

The main idea is that an experience fails to be veridical if it isn't *caused in the right way*. Like (A*), such a conception of veridical experience focuses on the connection between experience and the world—it presupposes that there is at least a causal relation between experiences and the world. But what does it mean to say that a veridical must be caused in the *right way*?

On this view, veridical experiences are the outcomes of well-functioning causal processes constitutive of various sensory channels, which allow subjects to be in direct contact with objects in their immediate environment. Such a conception of veridical experience relies essentially on a *functionalist* characterisation of veridical experience. The thought is that veridical experiences play a central *role*, mainly with respect to perception and perceptual knowledge. One cannot gain perceptual knowledge of certain objects and scenes without having veridical experiences of such objects and scenes. But experiences which are not causally connected to those objects and scenes in the right kind of way typically fail to fulfil this role.

¹⁷ In this sense, unlike (B), (A*) is compatible with a so-called 'purely *relational* conception' of experience—where this means that there is no intermediary between experiences and the world; in particular, no representational intermediary. For recent attempts to defend such a relational account of experience: see, e.g., Campbell (2002: ch. 6), Martin (2002), Maund (ch. 2), Travis (2004).

Suppose that a mad scientist uses electrodes attached to your brain to cause experiences of a kangaroo in front of you. Your experience results from a deviant causal chain. In such a situation, you are *hallucinating* a kangaroo, not perceiving one. One would not say that you are in a position to know the kangaroo on the basis of your experience. This is so, even if, by coincidence, there really is a kangaroo in front of you. Thus, according to this alternative conception of veridical experience, your experience isn't caused in the right way; and so it isn't veridical.¹⁸ Admittedly, it might be extremely complex to specify in detail what the 'right way' amounts to, or to distinguish precisely between normal and deviant causal chains. However, the partial success of causal theories of perception and causal theories of knowledge suggests that a similar account might be available for the notion of 'veridical experience'.¹⁹

This alternative account of veridical experiences might suit those theorists who think that, if there are limitations upon our discriminatory powers, they must be characterised in terms of (A) or (A*) *only*. The notion of 'veridical experience' is cashed out in causal terms only: as such, there is no need to think of perceptual experiences as representational. Hence, such a conception of veridical experience provides another explanation for the limitations on our ability to veridically experience something like a one nanometre growth.

Suppose that a change in size of just one nanometer is insufficient to bring about in the right kind of way any change in experience at all. As a result, such a change certainly cannot bring about an experience *of* a one-nanometer growth, and hence it is impossible for us to have a veridical experience as of a one-nanometer change in size. On such a view, changes in size which are smaller than ϵ cannot impact on a subject's experience *via* a normal causal chain: they are too small to trigger such effects. But note that this explanation for (A*) leaves open the possibility that such changes sometimes influence experience through a deviant

¹⁸ Compare Lewis (1980).

¹⁹ Note that things aren't any easier for the account of veridical experiences in terms of true representation that Graff takes for granted. Surely, we don't want to say that a creature that only has black and white vision cannot represent the world as it is, nor do we want to say that none of our perceptions are veridical due to the fact that they do not represent objects as collections of atoms. Typically, what counts as a correct representation of the world will depend on the kind of information that the subject's perceptual system is capable of conveying.

causal chain—in which case, illusions and hallucinations aren't ruled out by definition.

Graff could respond that this alternative account of veridical experiences fails to match all our intuitions about veridical experiences. But this misses the point. It does not matter here whether, in order to count as veridical, experiences must have a true representational content, or whether they must only be caused in the right way—or both. What is at issue is whether proponents of (A) can reformulate (A*) in such a way that (i) it is plausible, (ii) it does not collapse (A) into (B), and (iii) leads to non-transitivity. The causal account of veridical experiences suggests that what is required for such a reformulation is the distinction experiences caused in a deviant manner and experiences caused in the right way.

To be explicit, the causal account of veridical experiences suggests the following formulation of (A)—a formulation that captures the idea expressed in (A*), without collapsing it into (A**) or (B):

(A') For some small amount of change, whenever an object changes by any lesser amount, this change in the object cannot produce, in the right kind of way, any change in experience at all.

Note that (A') avoids the difficulties faced by (A). First, (A') does not entail that perceptual illusions of certain changes are impossible as (A) did. For instance, it is compatible with (A') that, whilst observing an object that is not changing, Winifred has a non-veridical experience—i.e., an experience produced in the wrong way—of a changing object. All (A') tells us is that such an experience must be illusory, and must have been caused in a somewhat deviant way. Second, (A') clearly does not collapse into (B). Like (A), (A') is explicitly concerned with the connection between the world and experience. Unlike (A**) and (B), (A') does not mention the notion of 'representation', nor does it attempt to specify the sort of experiences we can have.²⁰

²⁰ And again, note that (A**) and (A') are logically independent, so that one might be true without the other. Thus, unlike (A**), (A') admits the possibility that experiences *as of* one-nanometre changes are produced in illusory cases—i.e., when such experiences aren't produced in the right way. This might seem rather counter-intuitive, of course. But the fact that (A') is compatible with bizarre situations says nothing against it. To repeat, despite the logical independence of (A') and (A**), it is quite natural to regard both of them as true. Nevertheless, such a *de facto* connection between (A') and (A**) does not imply that (A') collapses into, or must be explained by,

Furthermore, (A') seems quite plausible. For example, watch a growing seedling for ten minutes and you will find that the minuscule change in the size of the plant does not have any effect on your experience. Rub your finger over the surface of a table and you will change the composition, texture and size of the table by removing atoms from the table and leaving behind atoms from your finger. Despite these small changes in the table, your experience of the table will, under normal conditions, remain fairly stable. That is, your perception of the table will remain unaffected by this small change. Indeed, every object around us is constantly undergoing minute changes in size, shape, and composition as molecules and energy are exchanged between the surfaces of objects and their environment. Such small changes do not seem to have any influence whatsoever upon our experiences.

2.6. *The Way to Non-transitivity*

All that remains to show is that (A') entails that *looks the same* is not transitive. Suppose that (A') is true and that Winifred observes an object for an interval of time during which that object changes by some sufficiently small amount for (A') to apply. The change in the object will be too small to produce any change in Winifred's experience. If Winifred's experience of the object does change, therefore, the change must have been brought about by a *deceptive* cause. That is, the change in Winifred's experience will not have been produced in the right way by a corresponding change in the object itself, but in some other way involving a deviant causal chain. And if a change is caused in such a way, we will say that the cause of the experience is *deceptive*. But in the absence of deceptive causes, Winifred's experience of the object will not change. And if her experience of the object does not change, the object must look the same throughout the period Winifred observes it.

If (A') is true,²¹ then, in the absence of deceptive causes, there will be instances of non-transitivity of *looks the same as*. That is to say, if (A') is true, the only way

(A**). Rather, the explanation could go in the other direction. Thus, our inability to perceptually represent certain very small changes—(A**)—is quite plausibly explained by the fact that such changes don't normally affect us perceptually—(A').

²¹ Together with the additional premises (P1) and (P2)—see footnote 14. However, we should note that (P1) is open to the same objections that Graff directs against (A). For (P1) states that whenever we are confronted with a

that *looks the same* could fail to be non-transitive would be for some deceptive cause to intervene on every occasion where we might apply (A') to derive an instance of (INT). Putting aside a radical global scepticism according to which *all* our perceptions are produced by deceptive causes, it would be very odd indeed if deceptive causes occurred in each and every situation where we might get an instance of non-transitivity. Certainly, the burden of proof is on the defender of (TRAN) to tell us why deceptive causes invariably pop-up in just the right places.

Wrapping Up: Graff is mistaken to claim that no reformulation of (A) can both avoid the problem of impossible illusory experiences, and provide a version of (A) that does not collapse into (B). We have shown that, at a crucial point in her argument for the collapse of (A**) into (B), Graff unwarrantedly replaces one possible reformulation of (A) with another. While (A**) might collapse into (B), we have argued that (A*)—or (A')—doesn't. Thus, we have offered a reformulation of (A) which avoids all the objections raised by Graff, and successfully entails that *looks the same as* is non-transitive.

3. Phenomenal Continua

The second consideration which, according to Graff, might support the non-transitivity of *looks the same as*, concerns the existence of *phenomenal continua*. By 'phenomenal continua', Graff means experiences of change in a given quality \mathcal{Q} (say, colour, shape, height, motion, etc.), which are experienced as *appearing continuous*. Think for instance of your experience of a white wall at sunset: if you look at the wall for the whole time, the wall might look to turn progressively and continuously orange. In this respect, phenomenal sorites series are indeed supposed to be instances of phenomenal continua.

change greater than ϵ we will perceive it as a change. But could it not be the case that an object that is actually changing is misperceived as not changing? (P1) can be modified so as to avoid making us infallible detectors of change:

(P1*) If an object in our perceptual environment changes with respect to quality \mathcal{Q} by an amount greater than ϵ , then the object will *usually* appear to us to have changed.

Admittedly, (P1*) is weaker than (P1). However, this weakening will not provide much solace to the defender of transitivity, since the purported examples of non-transitivity are cases where we do in fact correctly notice a change.

In the next section of her paper, Graff focuses on whether (i) there really are phenomenal continua, and (ii) whether their existence can provide a reason for the non-transitivity of *looks the same as*. She suggests that the following two-premise argument—the ‘Speculative Argument’ (Graff, 2001: 921)—might be advanced in support of (INT):²²

- (1) *There are phenomenal continua*: e.g. changes in colour across a spectrum, which look continuous.
- (2) *Small enough regions of phenomenal continua appear homogeneous*: there is a small enough width w which is such that all regions of width w look homogeneous.

Therefore,

- (3) *‘looks the same as’ is not transitive with respect to small enough regions of phenomenal continua*: adjacent regions, when they are sufficiently small, look the same, although the end regions don’t look the same.

Although the first premise is fairly transparent—after all, it is an argument exploiting the existence of phenomenal continua—, we cannot say the same about the second premise. Graff calls it the ‘Homogeneity Thesis’. The idea is that phenomenal continua can be segmented into small parts or regions, which look homogeneous: that is, they do not appear to change, not even continuously. Still, she doesn’t tell us exactly what work such a premise is supposed to be doing in the argument.²³

Although Graff concedes that the Speculative Argument is valid,²⁴ she resists it on various counts. Her general strategy is to establish that “we do not have

²² Graff claims that such an argument presents an informal version of Crispin Wright’s (1975) argument for the non-transitivity of *looks the same as*—an argument she discusses in the last section of her paper. We shall ignore Wright’s version of the argument, and restrict ourselves to her discussion of the informal version.

²³ In particular, it is unclear why the Homogeneity Thesis is needed at all. Presumably, there could be a phenomenal continuum composed of heterogeneous-looking patches—colour mosaics, say. Adjacent patches might look the same in their heterogeneity, and yet differ in such a minute way that non-adjacent patches will not look the same, but only similar. Premise (1) seems sufficient to entail (INT).

²⁴ Graff (2001: 922) argues that the Speculative Argument can be shown to be valid, given the additional principle that:

x looks homogeneous *if and only if* all sub-regions of x look the same as each other.

When applied to the second premise of the Speculative Argument, this principle entails that *looks the same as* is non-transitive.

sufficient grounds for accepting either premise” (2001: 922). Again, though, we find some of her objections unconvincing. Again, we shall defend the following conditional: if there are good reasons to believe in the existence of phenomenal continua and in the Homogeneity Thesis, such reasons are not seriously threatened by Graff’s objections.

First, we consider her attempt to undermine beliefs in the existence of phenomenal continua (§3.1). Graff then attempts to ‘explain away’ the Homogeneity Thesis, by advancing her own alternative explanation as to why there might seem to be (but aren’t really any) phenomenal continua (§3.2). We raise one worry about this explanation (§3.3), consider some possible responses on Graff’s behalf, and show why they fail (§3.4).

3.1. *Loose Talk*

A propos premise (i), Graff argues that we have no good reason to believe in the existence of phenomenal continua. As she rightly points out, we sometimes speak loosely for the sake of convenience. From this, she infers that, since we might sometimes speak loosely about the way things look, we shouldn’t take such talk as indication that certain things do look continuous. More generally, Graff argues, we cannot be sure that the colour spectrum does not in fact look discontinuous—that it doesn’t “really look to contain, at certain points, discrete but barely noticeable changes in colour” (2001: 922). And so, we have no good reason to believe in the existence of the kind of phenomenal continua that phenomenal sorites arguments typically exploit.

There is something puzzling about this sort of argument. Of course, it is true that (i) we sometimes talk without too much precision, and also that (ii) we are sometimes unsure whether things look a certain way rather than another. But the fact that we sometimes talk loosely for the sake of convenience—that the contents of our assertions lack precision—concerns the way we talk. How is this of any relevance to the way things look to us, or to our knowledge of the ways things look? When we do speak loosely for the sake of convenience, we are more than often aware of doing so. This means that we know very well that things don’t really look quite the way we say they look. Thus, the occurrence of loose talk should have little bearing on whether we can be sure that the chromatic variations on a spectrum actually *look* continuous.

However, Graff suggests that loose talk about phenomenal continua might be misleading in situations like the following:

[...] imagine we could have the following two experiences: the first, of a cursor on a computer screen looking to move discontinuously from one *pixel* to the next (imagine also that pixels are incredibly small), but in a very even way; the second of the cursor looking to move *continuously* from one pixel to the next. Again, the two experiences, although different, would not strike us as being different. We would describe both as experiences of continuous-looking motion.

We should not be misled into thinking that just because it may be convenient to describe a change as apparently continuous, that it really is that way. (Graff, 2001: 923)

The point of her argument seems to be that, in the situation just described, we are not *justified* by our experience to believe that the motion of one cursor on the computer screen looks continuous. Since we cannot discriminate the continuous-looking motion of one cursor from the discontinuous-looking motion of the other, we have no good reason to believe that the motion of one cursor looks continuous. And in particular, loose talk about the motion of these cursors isn't such a reason.

One difficulty with this point, though, owes to the fact that 'looks continuous', just like 'looks red' and 'looks the same as', is supposed to be taken in its phenomenal sense: presumably, it is a purely observational predicate too. And so, insofar as the cursor *appears* to move continuously to the perceiver, this should suffice to warrant such a perceiver in describing the cursor's motion as 'looking continuous'—especially if she cannot detect any difference between the two cursors. For all it takes for a perceiver to correctly apply this predicate in its phenomenal sense is that things do indeed *appear* continuous to her. At least, this is the way Graff described the conditions of application of such observational predicates (2001: 907).²⁵

Another worry concerns the scope of Graff's argument. Surely, from the fact that we can't be sure whether *some* things really look continuous rather than discontinuous, one can't infer that this uncertainty extends to all putative cases of

²⁵ We suspect that Graff is here oscillating between two different constraints on observational predicates. That is, whether (i) an appearance of continuous motion is sufficient for the object to look continuous, or whether (ii) it isn't sufficient, and the fact that the object actually moves continuously also determines whether its motion looks continuous. Or perhaps, Graff might be introducing yet another kind of predicate for uncertain phenomenal judgements, such as "appears to look continuous".

phenomenal continua. We might sometimes doubt whether chromatic variations on a spectrum from red to yellow look continuous. But sometimes we don't.²⁶ And all the proponent of the Speculative Argument really needs are just *some* instances of colour spectra, where all transitions from one patch to the other look continuous. Premise (1) in the Speculative Argument contains an existential quantification: it doesn't require that *all* colour spectra look continuous. Hence, it is unclear, at best, how Graff's point serves to undermine this first premise.

3.2. *Noticing the Way Things Look*

What about premise (2)—namely, that partitions of phenomenal continua of a certain size all look homogeneous? Having convincingly ruled out that the Homogeneity Thesis presents a *necessary* requirement for phenomenal continua,²⁷ Graff goes on to consider whether such a thesis might at least be contingently true, and “states an empirical fact about phenomenal continua as they are experienced by beings like us” (Graff, 2001: 925).

Here, Graff focuses on a particular example, which might seem to justify the claim that the Homogeneity Thesis can at least sometimes be true: experiences of slow motion.

In the case of position, the justification for taking as given that small enough changes are perceived by us as no change at all seems to be that sometimes objects move so slowly that it seems they look still—take the moon, or the hour-hand on a clock. The thought is that if a moving object seems to look still during an interval, then it must be because we cannot visually distinguish any of the positions it is during that interval—the *reason* it seems to look still is that our ability to visually discriminate does not extend to positions that are too close together. (Graff, 2001: 926)

²⁶ Compare the familiar point about perceptual illusions: we might sometimes doubt whether one of our experiences is veridical. But, *pace* the sceptic, this doesn't entail that we do—or even should—*always* doubt whether our experiences are veridical.

²⁷ She remarks that for non-phenomenal continua, no consideration of homogeneity is needed. And so, it is unclear why phenomenal continua, if they are real continua, should need the Homogeneity Thesis. This suggests that proponents of non-transitivity might indeed resort to a homogeneity-free version of the Speculative Argument.

She admits that the Homogeneity Thesis could serve to explain the occurrence of experiences of slow motion.²⁸ It might be that, in some cases, certain objects can look to change their position (or colour) in a continuous manner, *because they look not to change for short periods of time*. Over longer periods, however, their position looks to have changed. But there is no particular time at which they seem to be moving. And so, while the hour-hand looks the same with respect to its position after a twenty-seconds interval, and likewise for another similar interval, it may look to have a different position after a forty-seconds interval.

Yet, Graff insists, this is not the only available explanation. She goes on to advance an alternative account of what goes on in cases of apparent slow motion:

Another explanation, consistent with transitivity, is that when we look (say) at the hour-hand on a clock, although it does in fact look to change in position over the course of twenty-second intervals, the change in appearance is too slight, and too slow, for us to *notice* it. We *judge* the hour-hand to look still, but our judgement about the character of our experience is mistaken. (Graff, 2001: 927)

One might wonder how Graff's alternative explanation threatens premise (2) in the Speculative Argument. After all, the mere fact that there exists another putative explanation of cases of slow motion shouldn't by itself undermine the explanation grounded in the Homogeneity Thesis.

It is important to see that the strategy behind Graff's alternative explanation is not just to provide a re-description of cases of slow motion where the Homogeneity Thesis is false, but is more wide-ranging. Graff here attempts to 'explain away' the intuition behind the Homogeneity Thesis, by redescribing the

²⁸ However, she claims that the proposed description of the case makes experiences of constant motion impossible (2001: 927). The argument seems to proceed as follows: according to the proposed explanation, (1) the hour-hand looks still for intervals of twenty second because, during such intervals, the hour-hand moves for only $1/6^\circ$ on the clock, which isn't a perceptible difference. But then, (2) if changes of position of $1/6^\circ$ on the clock are too small to be perceived, the same should apply to the second-hand, which changes its location by that amount during $1/36$ second intervals. Thus, it seems, (3) the second-hand should look still too, if only for $1/36$ second intervals. The problem, says Graff, is that (4) the second-hand looks to be moving constantly, and thus never looks still. But why is this a problem? As she also acknowledges (*ibid.*), it is unlikely that anything looks any way during $1/36$ second intervals—at least for creatures like us. Our perceptual system needs time to process information. And so, the fact that the second-hand might have looked still for $1/36$ second intervals—had we been able to see anything during such intervals—seems compatible with the fact that the second-hand actually looks to move constantly in the course of longer intervals. Indeed, this seems to be the very phenomenon that cartoonists exploit, for instance.

phenomenological data allegedly explained by this thesis. According to her alternative explanation, it's not that objects look still for very short amounts of time. As a matter of fact, she claims, the hour-hand *looks* to move during twenty-second intervals. It's just that its apparent motion is too small to be *noticed* by us.

Thus, on her view, we wrongly describe the hour-hand of a clock as looking to have the same position during twenty-second intervals, because we cannot notice the way the hour-hand actually looks. Such a move has consequences, not just for the Homogeneity Thesis, but undermines the existence of phenomenal continua as well. In fact, Graff's alternative explanation generally challenges the availability of phenomenological descriptions of any example to which (INT) might seem to apply. It suggests that all such descriptions might be misguided, because we don't have any epistemic access to the phenomenology of such experiences.

Since so much rests upon it, Graff goes on to defend the distinction between (i) the *way things look* and (ii) our *noticing how they look*. She rightly observes that such a distinction is uncontroversial, and that even proponents of non-transitivity should accept it:

It should not be objected that the distinction introduced here—the distinction between how things look, and what we noticed about how they look—is an *ad hoc* distinction, tailored to the present purpose, since it is a distinction naturally invoked in other cases. Suppose a friend whom I see daily comes up to me and says, expectantly, 'Do I look different?' I may be unable to discover any change in her appearance. Still, once she goes on to tell me that she has lightened her hair (say), I may be able to tell that, yes, her hair colour does look different from how it did the day before. Yet it seems wrong to say that her hair colour looked different to me only *after* I was informed of the change. Instead, it looked different all along, but only after being informed did I *notice* that it looked different. My inability to discover the change, prior to being informed, was not a visual failing but an epistemic one. (Graff, 2001: 927-8)

So far, so good. We have one important reservation about Graff's alternative explanation, however. Her convincing defence of the distinction between (i) the *way things look* and (ii) our *noticing how they look* does not warrant her application of that distinction to cases of slow motion. Nor does it warrant, for the same reason, her strategy of phenomenological re-description.

3.3. *Unnoticeability & failures to notice*

Here is how the worry arises. On the one hand, it is possible that things we *fail* to notice can nonetheless look a certain way to us, as Graff's lightened-hair example makes plain. It seems very implausible indeed that the look of things—a phenomenological fact—be determined by whether or not we actually notice their look—an epistemological fact. In other words, appearances are independent from our epistemic access to them, at least in the following sense:

(N1) for some normal subject *S*, for some object *o* seen by *S*, and for some property ϕ ,
it is possible that (i) *o* looks ϕ (to *S*) and (ii) *S* does *not* notice that *o* looks ϕ .

According to (N1), an object can look a certain way, even though the subject does not *actually* notice the way that object looks. Note that (N1) does not rule out the possibility that the chromatic change in my friend's hair is *noticeable*.²⁹ One might occasionally fail to notice the change in one's friend's hair colour. Nevertheless, it is still the case that such a change *can* be noticed. For one thing, I *can* notice the difference, once my friend has attracted my attention to it. Yet, this would not be possible if the change was truly *unnoticeable*. For another, it seems that, even though I failed in this instance, I could have noticed the difference in hair colour without being told. After all, we often do notice such differences.³⁰

On the other hand, Graff comes close to acknowledging that differences in the position of the hour-hand over the course of twenty-second intervals *cannot* be noticed. As she puts it, such changes are “too slight, and too slow, for us to *notice* [them]” (*ibid.*). Indeed, it seems that we couldn't notice such changes *at all*, even if we looked very hard. Given the way the world is, and given the way our perceptual system functions, it seems impossible for us to (successfully) notice such

²⁹ By 'noticeable', we mean a *dispositional* property of the object, such that it is at least *possible* to notice it.

³⁰ Hence, Graff's claim that “I may be *unable* to discover any change in her appearance” (our emphasis) should read “I might actually fail to discover it”. Surely, Graff doesn't want to imply that one lacks the capacity to discover such things (otherwise, how could one notice the relevant difference after having been told?), or that it is impossible to notice them.

differences.³¹ Hence, in this case, the appropriate distinction between looks and our epistemic access to them should rather be along the lines of:

(N2) for some normal subject *S*, for some object *o* seen by *S*, and for some property ϕ , it is possible that (i) *o* looks ϕ (to *S*) and (ii) *S cannot* notice that *o* looks ϕ .

(N2), however, seems stronger than (N1), since it allows that more possibilities are ruled out. In particular, while (N2) demands that it is not possible for the perceiver to notice that *o* looks ϕ to her, (N1) doesn't.

Now, the problem with Graff's alternative explanation of cases of slow motion is that the distinction she needs in such cases between looks and our epistemic access to them hasn't been properly motivated. Although (N1) is well-supported by the lightened-hair example, the application of the distinction between (i) the *way things look* and (ii) our *noticing how they look* to cases of slow motion seems to require (N2). For it seems impossible that one notices that the hour-hand actually looks to be moving during short enough intervals. And it is not so clear whether the lightened-hair example supports the latter. As we have seen, (N1) is compatible with the *possibility* that one notices the change in colour in one's friend's hair. (N2), on the other hand, allows that such a possibility be ruled out.

So we agree with Graff that one ought to distinguish (i) the *way things look* from (ii) and the fact that we actually *notice the way things look*—in the sense specified by (N1). But the distinction Graff really needs is one between (i) the *way things look* and (iii) the *noticeability* of their looks—as in (N2). We conclude that Graff's alternative explanation is ill-motivated. The reason she offers for the distinction only supports the weaker (N1). And such a claim does not entail (N2).

3.4. *Phenomenal Looks once more*

There are at least three ways in which Graff might respond to the difficulty we raise for her alternative explanation of cases of slow motion. She might argue (i) that changes in hair colour are in fact *unnoticeable*—in which case, the lightened-

³¹ Of course, instances of 'illusory' noticing are always possible. Even if *o* is ϕ , it may be that the perceiver suffers from an illusion when she thinks that *o* looks ϕ . Since noticing seems to require success—and, importantly, to have been caused by a veridical experience *in the right way*—, such cases should not be taken as real instances of noticing.

hair example does indeed support (N2). Or she might insist that (ii) apparent changes of location of the hour-hand for twenty-second intervals *can* be noticed—so that (N1) applies to them too. Finally, she might point out that (iii) (N1) *does* apply to cases of slow motion anyway, no matter how we describe that case.

The first response seems implausible: it is unlikely that Graff will argue that the change in colour in one's friend's hair is *unnoticeable*—especially if she thinks that one can notice such a change, once one's attention has been drawn to it.

The second response is more interesting. Graff could retort that apparent changes in position of the hour-hand in the course of twenty-second intervals are in fact noticeable—in the sense that it is possible to notice them. Note that, ultimately, it is an empirical question whether or not it is possible for creatures with our perceptual faculties to notice such changes. In this respect, the question will not be decided by thinking of creatures with perceptual capacities radically different from ours. The sort of possibility under consideration here is psychological: it depends on the psychological faculties we actually have, and which we actually exercise. Thus, whether or not the motion of the hour-hand can be noticed during such intervals mainly depends on whether *we* actually notice it, at least sometimes.

In this respect, there is an important empirical difference between Graff's lightened-hair example and the case of the apparent continuous change of position of the hour-hand on a clock. In the former case, there are some readily available explanations as to why we sometimes fail to notice that our friend's hair colour has changed. It might be that (i) our attention is distracted and drawn to some other aspect of her appearance. Perhaps, (ii) the change in hair colour is not salient enough to attract our attention. Or (iii) there might be too many aspects of our friend's appearance to process, so that when she asks "Do I look different?", we're not quite sure where to look. Alternatively, (iv) it might be that we don't quite remember how our friend's hair colour previously looked. And without such information, it seems, we cannot notice any difference between the way it previously looked and the way it looks now. Such possibilities make perfectly good sense, and each one can help to explain why we don't notice that our friend's hair colour looks different.

But nothing of the sort is available with the motion of the hour-hand of a clock. We might look very hard at the hour-hand for very many twenty-second intervals without noticing any apparent motion. Insofar as we focus our attention on the hour-hand and try to see an apparent change in its location, no distraction of attention might occur, no reason why the subject's attention fails to be drawn towards the feature in question, no failure of memory.³² In the absence of any plausible explanation as to why we typically fail to notice the apparent motion of the hour-hand for twenty-second intervals, it seems reasonable to think that, in fact, we *cannot* notice such changes.³³

There is, however, a third response Graff might push against our claim that her use of the distinction between (i) the *way things look* and (ii) our *noticing how they look* in cases of slow motion is ill-founded. We pointed out that Graff's lightened-hair example supports (N₁), not (N₂). We then argued that *only* (N₂) is applicable to the case under consideration, because, unlike (N₁), (N₂) allows that that it is *not* possible to notice such changes. But Graff might answer that this is a mistake—that, in fact, *both* (N₁) and (N₂) apply to apparent changes of position of the hour-hand during twenty-second intervals, even if such changes are unnoticeable. On this ground, she might claim that (N₁) suffices to capture the

³² Of course, there could be such failures. But it seems implausible that, whenever we look at the hour-hand of a clock for more than twenty seconds, we always—*albeit* contingently—forget where it was located when we started looking at it.

³³ At this point, Graff could insist that changes of position of the hour-hand during twenty-second intervals are noticeable, at least in the sense that one could *acquire the capacity* to notice them. Perhaps, this is a capacity that experienced watchmakers possess, although normal observers don't. However, compare this suggestion with the claim that one can acquire the capacity to notice and distinguish the various facets or components of the taste of a wine. In the latter case, it seems, the novice wine-taster can be taught this new skill fairly easily, when guided by the professional wine-taster: namely, by having her attention drawn to this and that aspect of the taste of a particular wine. In this respect, those aspects of the taste are noticeable even prior to the acquisition of the new skill: they *can* be noticed just by focusing one's attention in the right way. Of course, the novice wine-taster might not have previously noticed them, but her newly acquired skill is not one that requires a profound modification of her perceptual capacities. On the other hand, it doesn't seem as though one could learn to notice apparent changes in the position of the hour-hand during twenty-second intervals in the same way. No matter how hard one tries to focus one's attention on the position of the hour-hand, it just doesn't look to move for such intervals. Thus, unless a substantial change occurs in one's perceptual apparatus, so that one can become sensitive to such slight changes of location, it doesn't seem as though such changes are at all noticeable.

distinction she needs for her alternative explanation, and that her appeal to such a distinction is perfectly well-motivated.

For, if (N1) allows that we *can* notice the way things look, it doesn't *entail* such a possibility. To claim that *not-P*—namely, that a subject does *not* notice the way things look—is neutral as to whether it is also possible, or not possible, that *P*. And so, Graff might argue, (N1) is also compatible with cases where it is impossible for us to notice how things look. In other words, (N1) applies just as well to cases of slow motion: strictly speaking, if we cannot notice such changes, it is trivially true that we don't notice them.

Even if this point is correct,³⁴ it fails to address the difficulty behind Graff's application of the distinction between (i) the *way things look* and (ii) our *noticing how they look* to cases of slow motion. The reason why (N1) is not applicable to such cases is that changes in the position of the hour-hand for twenty-second intervals do not *look* any way to normal perceivers. Note how both (N1) and (N2) express the independence of two kinds of conditions: in both claims, condition (i) concerns the way things look, while condition (ii) is about what we notice. Graff's response focuses on condition (ii), but ignores condition (i).

Thus, (N1) might *seem* to apply to apparent changes of position of the hour-hand during twenty-second intervals, *only because* it is trivially true that, being unnoticeable, such changes are not noticed. But in fact, (N1) does not apply in these cases, because it is unlikely that such minute changes in position can actually *look* to be moving during twenty-second intervals. Since such changes seem impossible to notice, it is unclear what ground there could be for the claim that they look any way to us. Admittedly, (N2) could offer such a ground, but it seems plainly false—at best, unwarranted.

Here is why. The sense of 'look' under consideration in claims like (N1) and (N2)—and indeed, in this whole discussion—is phenomenal. Thus, it becomes extremely puzzling how a change in the world could have a direct impact on the phenomenology of a subject's experience, although it is impossible for the subject to become aware—i.e., notice—of such impact. Of course, we are not always aware of all aspects of our ongoing phenomenal life. And there are many such aspects we

³⁴ Indeed, if we have shown that (N1) and (N2) are distinct, and that only the former—but not the latter—applies to the lightened-hair case, it doesn't follow that only (N2)—but not (N1)—applies to cases of slow motion.

often fail to notice. And perhaps, there are some properties of the phenomenology of our experiences that are utterly inaccessible to us. Nevertheless, it seems highly plausible to think that, if something does contribute to the phenomenal life of a subject, such a thing is at least *accessible* to her, so that it is possible for her to notice it. Thus, if the sense of ‘look’ under consideration is the phenomenal sense, we take the following assumption to provide another very plausible constraint upon the use of such a term:

(LN) for any normal subject *S*, and for any object *o* seen by *S*, *o* looks ϕ (to *S*) *only if* it is possible for *S* to notice that *o* looks ϕ .

Indeed, we shall consider (LN) as expressing the default position regarding the phenomenal sense of ‘look: unless a good argument can be presented against it, it seems to capture one crucial aspect of what we mean by the use of such a term.

Now, Graff’s claim that the hour-hand of a clock looks to move—in the phenomenal sense—in the course of twenty-second intervals, together with her own admission that such changes are too small to be noticed, violates assumption (LN). Given the size of the changes in question, it’s not even possible to notice them. In which case, it is highly unlikely that the hour-hand looks to move during such intervals. If this is correct, Graff has to show, either (i) that the hour-hand’s apparent change of location for twenty-second intervals can be noticed after all, or (ii) that (LN) does not constrain the phenomenal sense of ‘look’. As we have seen, the first option seems unlikely: at best, it would require more argument. The second option seems even more remote. To repeat, if an object is seen by a normal perceiver and looks a certain way—in the phenomenal sense, this should make a difference to the phenomenology of that subject’s experience, in the sense that it is at least possible for her to notice it.³⁵ Anything short of satisfying (LN), we suggest, simply doesn’t contribute to the phenomenology of experience.

³⁵ Again, we seem to detect a certain tension in Graff’s attitude regarding the phenomenal sense of observational predicates—see also footnote 25. It is worth pointing out that Graff’s motivation for the distinction between (i) the way things look and (ii) our noticing how they look seems to presuppose something like (LN). Why should one think that the change in one’s friend’s hair colour looks any way to one, if one fails to notice such a change? Assumption (LN) seems to provide the reason. Together with the fact that (i) one can later notice the difference, or that (ii) one could have noticed it, (LN) suggests that the change in one’s friend’s hair colour does indeed look some way to one, even if it isn’t actually noticed. Thus, it looks as though Graff is in no position to reject (LN). If

To summarise: we have tried to show that Graff's attempts to undermine the plausibility of each premise of the Speculative Argument wear rather thin. Although she argues that we might be mistaken to think of certain cases as phenomenal continua, it is unclear whether her argument succeeds—given the phenomenal sense of 'looks continuous'. But even if it did succeed, it falls far short of establishing that there are no phenomenal continua at all. As for Graff's attack on the Homogeneity Thesis, her alternative explanation of the phenomena under consideration seems to be plagued by more difficulties than the Homogeneity Thesis itself.

4. Conclusion

So what to make of this fairly complex exchange of arguments and counter-arguments for and against (INT)? In her paper, Graff intends to shift the burden of proof onto the advocates of (INT), by insisting that none of the familiar considerations advanced in favour of such a thesis actually succeed. Although we welcome Graff's attempt to resurrect the debate, we have argued that most of her objections fail to undermine either the argument from discriminatory limitations, or the argument from phenomenal continua—the 'Speculative Argument', as she calls it.

We have shown that the existence of limitations upon our perceptual capacity to discriminate in the sense captured by (A) can be formulated in a way that avoids all the difficulties raised by Graff. Such a formulation—(A')—can then be used to entail (INT). We have also shown that Graff doesn't quite succeed in throwing doubt on the existence of phenomenal continua. This means that the first premise of the Speculative Argument is fairly secure, but also that discriminatory limitations in the sense captured by (B) can also serve to support the claim that *looks the same as* is non-transitive.

At this point, one might think that we have two cogent arguments for (INT), and that even one such argument would be enough. One must keep in mind, however, Graff's suggestion that certain things may look to change without us noticing, since it has the potential to undermine any putative example of non-

she did, she would undercut a major motivation for her distinction between (i) the way things look and (ii) our noticing how they look.

transitivity. If Graff's alternative explanation did work, it would provide a basis for re-describing any situation where two objects appear to look the same as one where, in fact, such objects look different—it's just that we don't notice the difference. As a result, (INT) would remain uninstantiated. We have suggested, however, that such a strategy undermines the phenomenal sense of 'look', which was supposed to be granted for the purpose of this discussion. In this respect, her alternative explanation of the examples under considerations simply changes the subject, since she seems to rely on a sense of 'looks' which doesn't respect the constraints on its phenomenal sense.

These last remarks suggest, in the guise of a conclusion, a slightly different way to motivate the thought that *looks the same as* is not transitive. Assumption (LN) in particular can provide a way to derive the non-transitivity of *looks the same as*. Of course, we have not attempted to defend (LN). Rather, we take it to specify the default position about the phenomenal sense in which things might be said to 'look' a certain way. Suffice it to note that Graff is in no position to reject (LN), since her argument for distinction between (i) the *way things look* and (ii) our *noticing how they look* seems to presuppose it.

Here is the argument:

- (1) if a subject *S cannot* notice any change in the position of the hour-hand of a clock in the course of a twenty-second interval from time t_1 to t_{20} , then the hour-hand doesn't look to move between t_1 and t_{20} .
- (2) if a subject *S cannot* notice any change in the position of the hour-hand of a clock in the course of a twenty-second interval from time t_{20} to t_{40} , then the hour-hand doesn't look to move between t_{20} and t_{40} .
- (3) if a subject *S can* notice a change in the position of the hour-hand of a clock in the course of a forty-second interval from time t_1 to t_{40} , then the hour-hand does look to move between t_1 and t_{40} .
- (4) If the hour-hand doesn't look to move between t_1 and t_{20} , and doesn't look to move either between t_{20} and t_{40} , but it does look to move between t_1 and t_{40} , then *looks to be in the same position as* is not transitive.
- (5) for most normal subjects *S*, *S cannot* notice any change in the position of the hour-hand of a clock for twenty-second intervals between, say, t_1 and t_{20} , as

well as between t_{20} and t_{40} , but S can notice changes in the position of the hour-hand for forty-second intervals, say, between t_1 and t_{40} .

Therefore,

(6) *looks the same as* is not transitive with respect to location.

Each of the first three premises relies on (LN)—or its contrapositive, to be precise. Premise (4) conjoins the consequents of these three premises to derive the non-transitivity of *looks the same as*. So far, however, the premises only contain conditional claims.

Premise (5) advances an empirical claim about what is noticeable. Perhaps, such a claim might turn out to be false: perhaps, normal subjects can only notice changes in the position of the hour-hand of a clock in the course of *sixty* second intervals. In this case, together with the necessary adjustments, the argument will require an extra-conditional premise at the beginning. Even so, the accurate empirical claim, together with the relevant set of conditionals, entails that *looks the same as* is non-transitive, by re-iterated applications of *modus ponens*.³⁶

It is worth insisting that the conclusion of this argument supports only the non-transitivity of *looks the same as*. It may not suffice to justify the existence of phenomenal sorites series, as required by phenomenal sorites arguments. For all we have argued so far, it might well be that for at least one twenty-second interval, the hour-hand on most clocks does look to move discretely. In other words, we have only defended the claim that *some*—i.e., at least three—adjacent parts of phenomenal continua can look the same as each other. Perhaps, these are the only phenomenal continua that exist, and there are no phenomenal continua big enough to generate phenomenal sorites arguments.³⁷

³⁶ A similar argument can be constructed in the case of colour. Suppose that normal subject cannot notice any chromatic difference between patch #1 and patch #2, and likewise for patch #2 and patch #3, but that they can notice a chromatic difference between patch #1 and patch #3 (or patch #4). This fact, together with re-iterated instances of the contrapositive of (LN), entails that *looks the same as* is not transitive with respect to colour.

³⁷ *Acknowledgements:* A version of this paper was presented at the ANU Philosophy Society in Canberra in November 2003. Particular thanks go to Brett Calcott, Frank Jackson, Laurie Paul, Denis Robinson, and Kim Sterelny, for their incisive questions.

References:

- BARNES, Jonathan. 1982. 'Medicine, Experience and Logic', in Barnes, J., Brunschwig, J., Burnyeat, M., and Schofield, M., eds., *Science and Speculation* (Cambridge: Cambridge University Press).
- BURNS, Linda Claire. 1991. *Vagueness: An Investigation into Natural Languages and the Sorites Paradox* (Dordrecht: Kluwer).
- CAMPBELL, John. 2002. *Reference and Consciousness* (Oxford: Oxford University Press).
- CLARK, Austen. 1993. *Sensory Qualities* (Oxford: Oxford University Press).
- GOODMAN, Nelson. 1951. *The Structure of Appearance* (Cambridge, MA.: Harvard University Press).
- GRAFF, Delia. 2001. 'Phenomenal Continua and the Sorites', *Mind*, 110: 905-35.
- GREENOUGH, Patrick. 2003. 'Vagueness: A Minimal Theory', *Mind*, 112: 235-81.
- HARDIN, C.L. 1988. 'Phenomenal Color and Sorites', *Noûs*, 22: 213-34.
- HYDE, Dominic. 2002. 'Sorites Paradox', *The Stanford Encyclopedia of Philosophy*.
- JACKSON, Frank. 1977. *Perception* (Cambridge: Cambridge University Press).
- JACKSON, Frank, and PINKERTON, R.J. 1973. 'On an Argument against Sensory Items', *Mind*, 82: 269-72.
- LEWIS, David. 1980. 'Veridical Hallucination and Prosthetic Vision', *Australasian Journal of Philosophy*, 58: 239-49.
- MARTIN, M.G.F. 2002. 'The Transparency of Experience', *Mind and Language*, 17(4): 376-425.
- MAUND, Barry. 2003. *Perception* (Montreal: McGill-Queen's University Press).
- MERRICKS, Trenton. 2001. 'Varieties of Vagueness', *Philosophy and Phenomenological Research*, 62: 145-57.
- MILLS, Eugene. 2002. 'Fallibility and the Phenomenal Sorites', *Noûs*, 36(3): 384-407.
- RAFFMAN, Diana. 2000. 'Is Perceptual Indiscriminability Nontransitive?', *Philosophical Topics*, 28(1): 153-75.
- SAINSBURY, Mark, & WILLIAMSON, Timothy. 1997. 'Sorites', in Hale, B. & Wright, C., eds., *A Companion to the Philosophy of Language* (Oxford: Blackwell).
- READ, Stephen. 1994. *Thinking about Logic* (Oxford: Oxford University Press).
- TRAVIS, Charles. 2004. 'The Silence of the Senses', *Mind*, 113: 58-94.
- WILLIAMSON, Timothy. 1990. *Identity and Discrimination* (Oxford: Blackwell).
- WILLIAMSON, Timothy. 1994. *Vagueness* (London: Routledge).
- WRIGHT, Crispin. 1975. 'On the Coherence of Vague Predicates', *Synthese*, 30: 325-65.