Noise and Perceptual Indiscriminability

Benj Hellie

Perception represents colours inexactly. This inexactness results from phenomenally manifest noise, and results in apparent violations of the transitivity of perceptual indiscriminability. Whether these violations are genuine depends on what is meant by 'transitivity of perceptual indiscriminability'.

Poor Nelson! He wants to select *exactly* the right shade of paint for his dining room walls. Unfortunately, at the paint store, he is not up to the task:

Nelson's Predicament

Nelson is confronted with three paint chips, a', a'', and a'''. Nelson undergoes three successive visual experiences (particular events of visually experiencing), e_1, e_2, and e_3, in which he compares the colours of the chips (maximally determinate colour types or repeatables instanced in the chips): in e_1, he compares a' and a''; in e_2, a'' and a'''; and in e_3, a' and a'''.

Nelson's visual experiences are nonaccidentally veridical; he has a reasonable amount of time to compare the seen objects; the lighting is bright, colourless, and spatially and temporally invariant; the colours of a', a'', and a''' are spatially and temporally invariant; being paint chips, the textures of a', a'', and a''' are flat—in short, conditions are *optimal* for perceiving colours—or at least as close to optimal as mortals can ever attain. Moreover, the colours of other things are behaving unexceptionally; Nelson is not 'spectrally inverted', is not becoming spectrally inverted, and is in no danger of becoming spectrally inverted; Nelson has no misleading opinions about any of this, and is not in danger of acquiring any; no sceptics are lurking around Nelson, or around us, to threaten to destroy knowledge. In short, conditions are *normal*. Finally, all that Nelson has to go on in his experience for distinguishing colours are the ways the colours look 'in themselves' in the experiences: his experiences include no trusted friends informing him whether colours are different, no special instruments which flash

Mind, Vol. 114 . 455 . July 2005
doi:10.1093/mind/fzi481

> green lights when colours are the same and red lights when they are different. In short, Nelson's experiences are *unaided*.
>
> Despite this, something unfortunate for Nelson's purposes happens: while the colours of a' and a'' are indiscriminable on the basis of e_1, and the colours of a'' and a''' are indiscriminable on the basis of e_2, the colours of a' and a''' are discriminable on the basis of e_3.

Clearly the colours of a' and a''' differ—Nelson has discriminated them—but is the colour of a'' different from that of both a' and a''', or the same as that of one; and if so, which? Nelson doesn't know; he concludes that he can't determine exactly which colour something has just by looking.

Section 1 of this paper will explain what I believe is going on in Nelson's Predicament: a source of the phenomena is the inexactness of perception of colours, which is inescapable due to 'noise' in the perceptual signal path.[1] Section 2 discusses a competing treatment of Nelson's Predicament according to which the phenomena stem from the shiftiness, rather than the inexactness, of colour perception. The competing treatment is opposed to a traditional view stemming from Goodman (1951) on which perceptual indiscriminability is nontransitive. My conclusion on the status of the traditional view will be equivocal: on one way of understanding it, it is correct; on another, it might be false. Section 3 briefly evaluates some conclusions that have been drawn from cases like Nelson's Predicament.

A central theme in the paper is that overlooking the inexactness of colour perception has prevented an adequate treatment of indiscriminability; in particular, it has contributed to a tendency to conflate indiscriminability of two things with their being represented in exactly the same way.

1. Noise as a ground of nontransitivity

1.1 Nontransitivity and inexact perceptual representation

In this subsection, I argue that if perceptual representation of colours is inexact, a case like Nelson's Predicament can support the nontransitivity of perceptual indiscriminability. I begin by explaining indiscriminability on the basis of a perceptual experience in general. I characterize the sense in which perception of a colour can be inexact, and argue that

[1] The other philosophical treatment of perceptual noise of which I know is Hardin 1988, pp. 171–88.

if perception of two colours is inexact in this sense, they can be indiscriminable without being perceived alike. I then argue that if this can happen, it might be that perceptual indiscriminability is nontransitive (in either of two equivocal senses). In the following subsection, I explain why perception of a colour will, in general, be inexact.

1.1.1 Perceptual indiscriminability in general

What is it for two colours (or, for that matter, two things of whatever sort) to be indiscriminable on the basis of a perceptual experience—henceforth, sometimes just an *experience*—*e*? For two things to be indiscriminable is for it to be impossible to discriminate them; and to discriminate *a* from *b*, Williamson (1990) has convincingly argued, is to acquire knowledge of, or *recognize*, the distinctness of *a* from *b*. For *a* and *b* to be indiscriminable on the basis of *e* is for it to be *restrictedly* impossible to discriminate them—impossible to recognize their distinctness without appeal to resources beyond *e*.[2]

What are the resources an experience provides for recognition? Plausibly, recognition that *p* requires the incompatibility of not-*p* with one's evidence (cf. Lewis 1996). Experiences provide evidence; so if it is impossible to recognize the distinctness of *a* and *b* without appeal to resources beyond *e*, this would be due to the compatibility of $a = b$ with the evidence *e* provides.[3]

How does an experience provide evidence? According to a central tradition, experiences have *representational contents*—represent that the world and things in it are a certain way, which things may or may not be—and can therefore be assessed for *accuracy* or *inaccuracy* (Evans 1982; Harman 1990; Pryor 2000). (For ease of expression, I will adopt the assumptions of this tradition; I hope that those friendly to other traditions will be able to adapt my treatment.) Unless the content of one's experience is inaccurate if not-*p*—or, as I shall sometimes put it, unless one's experience *represents that p*—one's experience does not provide evidence against not-*p*, and one's experience cannot promote a judgement that *p* to the status of a recognition.[4] Conversely, if the content of one's experience is inaccurate if not-*p*, then one's experience

[2] My approach to indiscriminability and sorites puzzles is similar in other significant respects to Williamson's, especially in its emphasis on inexact mental representation. For a significant disagreement, see fn. 21.

[3] I use the notion of evidence in a way hopefully neutral with respect to Williamson's thesis that knowledge is evidence (Williamson 1995), since the focus is on the *provision* of evidence, rather than on the evidence itself.

[4] Perhaps other aspects of an experience—its 'quale', for instance—could be involved in promoting a judgement to a recognition; I will ignore this possibility.

does provide evidence against not-p, and could promote a judgement that p to the status of a recognition if the circumstances of the judgements are of the right sort (I will assume that the circumstances of Nelson's judgements are of the right sort). If it is impossible to recognize the distinctness of a and b, on this tradition, this would be because the content of one's experience is not inaccurate if $a = b$.

Of particular value in warranting a judgement of the distinctness of a from b would be content to the effect that a has some feature that b lacks. It would then be incompatible with the accuracy of the content of one's experience that a is b.

1.1.2 Perceptually indiscriminable colours

Plausibly, the content of experience includes information concerning not only which objects have which properties, but also which properties have which *higher-order* properties; it seems that whether an experience provides evidence for the distinctness of two colours is determined by which higher-order properties it attributes to them. For instance, an experience e^* of red, orange, and green objects r, o, and g represents, and is inaccurate if it is not the case, that redness, orangeness, and greenness stand in a certain ternary relation of comparative similarity, with red and orange more similar than either is to green; it also represents that redness and orangeness have the monadic higher-order property of being 'warm', while greenness has the monadic higher-order property of being 'cool' (and therefore not warm). The content of e^* is therefore incompatible with the identity of the colour of r with the colour of g, and could therefore warrant an appropriately situated judgement of the distinctness of the colour of r and the colour of g. These colours are thus discriminable on the basis of e^* (if external conditions are right). Conversely, unless the content of e includes the representation of the colour of a as having some higher-order property that the colour of b lacks, the content of e would be correct even if the colour of a and the colour of b were identical; in this case, these colours would be indiscriminable on the basis of e. The alternative would be that the non-identity of these colours is just a *primitive part* of the content of e, not grounded in any represented difference in their higher-order properties. But this would seem to require the possibility that two colours could look to be different without there being any difference in the way they look, and I have a very hard time taking seriously the possibility that there could be such a primitive ungrounded sensation of difference in colour.

One sort of higher-order property an experience represents a colour as having is its *location in the colour solid*.[5] For instance, the specific shade of red *r* is represented as having by e^* is represented as being located among the warm colours, as being located among the highly saturated colours, as being located among the reds, as being located among the scarlets, as being located among the darker scarlets .... (Other higher-order 'internal' properties of a colour, such as the comparative similarity relations in which it stands, its 'warmness/coolness' status, saturation, brightness, and the like are all determined by its location in the colour solid.) *e* would then be incompatible with the identity of the colour of *a* and the colour of *b*, and provide evidence against this identity, if it is incompatible with *e*'s accuracy that the colour of *a* is outside a certain region R_a or that the colour of *b* is outside a certain region R_b, where the two regions do not overlap. Under the right circumstances, a judgement of the distinctness of these colours could be warranted by *e*. However, if R_a and R_b overlap, then it is compatible with *this* aspect of the content of *e* that the colour of *a* and the colour of *b* are both in the region of overlap; moreover, that the colours are identical. This aspect of the content of *e* would not count as evidence against the distinctness of the colours. If *e* affords no other evidence against their distinctness—*e* is unaided—the colours would be indiscriminable on the basis of *e*.

On this analysis, there can be indiscriminability without sameness of representation. R_a and R_b need merely *overlap*; they do not need to be identical.

1.1.3 Nelson's Predicament, inexactness, and shiftiness

In the rarefied atmosphere of the paint shop, experience is unaided; the higher-order properties experience attributes to a colour by which it might be distinguished from others seem limited to those determined by ostensible locations in the colour solid. The following explanation of Nelson's Predicament becomes available.

Let $R(a, e)$ be the most specific region of colour space that the colour of *a* is represented as being located in by *e*. Then the content of e_1 is incorrect if either the colour of a' is outside $R(a', e_1)$ or the colour of a'' is outside $R(a'', e_1)$. Since a' and a'' are indiscriminable on the basis of e_1, these regions overlap. Since nothing else in the content of

[5] For convenience, I make some assumptions about the structure of the colour solid: it is structured like a subregion of the solid composed of triples of real numbers; maximally determinate colours, of the sort instanced by particular monochromatic surfaces, are points of the colour solid (so that the colour solid is not 'gunky'). All my treatment requires is that the region a colour is perceived as being located in is of significantly greater extent than the region it in fact occupies.

e_1 distinguishes these colours, the colours are indiscriminable on the basis of e_1. The content of e_2 is incorrect if either the colour of a'' is outside $R(a'', e_2)$ or the colour of a''' is outside $R(a''', e_2)$. Since a'' and a''' are indiscriminable on the basis of e_2, these regions overlap. Since nothing else in the content of e_2 distinguishes these colours, the colours are indiscriminable on the basis of e_2. The content of e_3 is incorrect if either the colour of a' is outside $R(a', e_3)$ or the colour of a''' is outside $R(a''', e_3)$. But $R_3{}'$ and $R_3{}'''$ don't overlap; hence (since e_3 is normal and optimal) the colours are discriminable on the basis of e_3.

Are there regions of the colour solid R', R'', and R''' such that $R' = R(a', e_1) = R(a', e_3)$, $R'' = R(a'', e_1) = R(a'', e_2)$, and $R''' = R(a''', e_2) = R(a''', e_3)$: that is, could it be the case that Nelson's total course of experience throughout Nelson's Predicament is *constant*—that is, that in Nelson's Predicament the region a colour is represented as being in (or as being not outside of) does not vary when it is presented in an experience and then presented in a later experience? If so, R' and R'' overlap, as do R'' and R''', but R' and R''' don't. If all these regions are of positive (greater than zero) extent, this is a mathematical possibility; otherwise, not.

What would it mean for these regions to be of positive extent? Recall that in my discussion of representation of colours as in locations in the colour solid, I wrote that the specific shade of red r is represented as having by $e^\star$ is represented as being located among the warm colours, as being located among the highly saturated colours, as being located among the reds, as being located among the scarlets, as being located among the darker scarlets …. What it would mean for the regions to be of positive extent is that (for Nelson, at least) the ellipsis does not bottom out with a maximally specific region containing only a single colour.

Must we take Nelson's experiences to represent the colour of each chip as being located in a maximally specific region of colour space? This can be taken in two ways, depending on the relative scopes of 'represent' and 'a maximally specific region'. If 'represent' is given wide scope (e.g. e_1 is incorrect if: there is no maximally specific region R such that: the colour of a' is in R), then the answer seems to be 'yes': under normal circumstances, every colour looks to be maximally determinate, no colour looks to be vague. If something looks orange, it looks to have a maximally specific shade of orange. But if 'represent' is given narrow scope (e.g. there is some maximally specific region R such that: e_1 is incorrect if: the colour of a' is outside R), the answer is less clear. If on the latter reading the answer to the question is no—so that it is not the

case that there is some maximally specific region R such that e represents a certain colour as being in R—I will say that e represents the colour *inexactly*. The thesis that all normal perceptual representation of colour is inexact can be stated as follows: for all a, for all (possible, normal, optimal, unaided) e, $R(a, e)$ has a positive radius.

If Nelson's experience is exact, it is *shifty*, not constant. Jackson and Pinkerton (1973) pioneered a treatment of cases like Nelson's Predicament on which experiences are exact and shifty; I will discuss this treatment in detail in section 2.2, below. Shiftiness and inexactness are compatible: a course of experiences might be both shifty and inexact. While I think that every medically possible course of experiences is in fact both shifty and inexact, I will focus the discussion for the time being by treating Nelson's Predicament solely by appeal to inexactness. In section 2.2, I will assess the consequences of mixing shiftiness with inexactness.

At this point, I am in a position to argue against exactness. At least one of the following inequalities holds: $R(a', e_1) \neq R(a', e_3)$, $R(a'', e_1) \neq R(a'', e_2)$, or $R(a''', e_2) \neq R(a''', e_3)$; suppose the first. Supposing exactness, $R(a', e_1) \cap R(a', e_3)$ is empty. But then Nelson's Predicament is impossible. Among the stipulations of the case are that Nelson's experiences are veridical—so that the colour of a' is, in fact, as it is represented as being, in both $R(a', e_1)$ and $R(a', e_3)$. But this is possible only if the colour of a' changes between e_1 and e_3—and among the stipulations of the case are that the colours of the objects he sees are unchanging. To the extent that Nelson's Predicament strikes one as possible, this should undermine the support for exactness. When we find ourselves in such a case in the real world, it is more attractive for one who wishes to preserve exactness to assume slight nonveridicality than it is to assume slight changes in colours; and to the extent that one thinks that Nelson's Predicament captures something pervasive about our condition, one must then accept pervasive slight nonveridicality. In my book, this cuts against exactness; in the next subsection I will sharpen this line.

1.1.4 What is the 'nontransitivity of perceptual indiscriminability'?

I have so far discussed the indiscriminability of two colours on the basis of a particular experience. This does not bear without further discussion on the transitivity of perceptual indiscriminability in general: the particularity needs to be generalized away. There are two strategies for doing this.

The first strategy quantifies the *e* position in '*a* and *b* are indiscriminable on the basis of *e*': *a* and *b* are perceptually indiscriminable *tout court* $=_{df}$ *a* and *b* are indiscriminable on the basis of any possible unaided experience occurring under normal, optimal conditions.[6] If it is assumed that what goes with respect to the discriminability status of *a* and *b* for some such experience goes for all (I will assess this assumption in section 2.2),[7] Nelson's Predicament shows that this notion of perceptual indiscriminability *tout court* is nontransitive.[8]

A second strategy factors the indiscriminability of the colours of *a* and *b* on the basis of *e* into a function from possible, normal, optimal, unaided experiences and objects into abstracta, and a relation between abstracta, and assesses this relation for its transitivity status. Compare: the relation of being taller than, holding between people at times, factors into a function from people and times to abstracta (the heights), and a relation between these abstracta (the > relation); Tim is taller than Michael in year y_1 just in case $H(t, y_1) > H(m, y_1)$ (Tim's height in y_1 exceeds Michael's height in y_1). Similarly, this second strategy for characterizing a general notion of perceptual indiscriminability requires a function $E(a, e)$ from an object and an experience into an abstract entity and a relation $\sim$ between abstracta such that *a*'s and *b*'s colours are indiscriminable on the basis of *e* just in case $E(a, e) \sim E(b, e)$. The question is then whether $\sim$—the *indiscriminability-making relation*—is transitive. A precondition for this style of factorization is the existence of a sort of abstract entity such that whether the colours of two objects are discriminable on the basis of a possible, normal, opti-

[6] '*a*' and '*b*' being read rigidly. A *metameric* pair are two surfaces which reflect light that the eye treats equivalently under normal circumstances—when the incident light is white—but which reflect light that the eye treats differently under abnormal circumstances—for example, when the incident light is pure spectral green. The colours of a metameric pair are perceptually indiscriminable *tout court* in the sense I have defined while being discriminable under abnormal conditions. I don't think this violates linguistic intuition, since generic predications are typically understood with reference to normal conditions: we say that John is a happy fellow even if after a death he would be unhappy; more to the point, we say that identical twins Tweedledee and Tweedledum are indiscriminable even if there is a difference in their weights detectable only by a very precise scale.

[7] Maybe the notion of optimal conditions admits conditions of too many sorts to licence this. In that case, no harm would be done to the project of demonstrating the influence of noise on nontransitivity to fix on a precisely specified set of external circumstances, since some noise is due to internal, biological features.

[8] The colour of a' and the colour of a'' are indiscriminable on the basis of e_1, hence by the assumption perceptually indiscriminable *tout court*; the colour of a'' and the colour of a''' are indiscriminable on the basis of e_2, hence by the assumption perceptually indiscriminable *tout court*; the colour of a' and the colour of a''' are discriminable on the basis of e_3, hence not perceptually indiscriminable *tout court*.

mal, unaided experience is wholly determined by how those objects stand to abstracta of that sort with respect to the experience.

The strategy in line with my approach takes the abstracta to be regions of colour space and ~ to be the relation of overlapping. The data of Nelson's Predicament are that $R(a', e_1) \sim R(a'', e_1)$, $R(a'', e_2) \sim R(a''', e_2)$, and $\neg(R(a', e_3) \sim R(a''', e_3))$. These data would establish the nontransitivity of ~ in the presence of $R(a', e_1) \sim R(a', e_3)$, $R(a'', e_1) \sim R(a'', e_2)$, and $R(a''', e_2) \sim R(a''', e_3)$: the first of these is equivalent to the claim that, fixing the region the colour of a' is represented as located in by e_1 and the region the colour of a' is represented as located in by e_3, for a possible, normal, optimal, unaided experience in which two colours were represented as being in these regions, the colours would be indiscriminable on the basis of that experience. I will assess the plausibility of these additional assumptions in section 2.2, below.

The claim that *a* and *b* are indiscriminable *tout court* factors into $(\forall e)R(a, e) \sim R(b, e)$. Nontransitivity of indiscriminability *tout court* thus requires nontransitivity of the indiscriminability-making relation, but not the other way around. When I discuss the auxiliary hypotheses in section 2.2, I will conclude that while the indiscriminability-making relation is nontransitive, indiscriminability *tout court* probably is not. Whether 'perceptual indiscriminability is nontransitive' is therefore equivocal.

1.2 Noise and inexactness

I now turn to defending the thesis that all normal colour perception is inexact. I will provide two arguments stemming from the unavoidable presence of noise in perception, one metaphysical and one phenomenological.[9]

1.2.1 An argument from the metaphysics of content

The properties represented in colour experiences are maximally determinate colours, and yet they are not represented maximally determinately.[10] I will argue that no plausible naturalistic theory of the

[9] I thus endorse the claim that our powers of discrimination are finite (cf. Wright 1975, p. 346). Graff (2001, pp. 916–7) wonders (paraphrasing somewhat): does this mean that there is some small amount δ such that (a) if one sees *a* and *b* which *in fact* differ in colour by less than δ, it will look to one as if the colour of *a* = the colour of *b*; or that (b) if one sees *a* and *b* which *look to one* to differ in colour by less than δ, it will look to one as if the colour of *a* = the colour of *b*? She points out that (a) is incompatible with misperception of differences less than δ, while (b) doesn't yield nontransitivity. But a third option is that (c) if one sees *a* and *b* which look to one to differ in colour by less than δ it will neither look to one as if the colour of *a* = the colour of *b* or as if the colour of *a* ≠ the colour of *b*. There is never apparent exact sameness, only absence of apparent difference. In practice, we often treat such cases as involving exact sameness, but this involves a step beyond what is revealed in perception.

metaphysics of content would allow for maximally determinate colours to be represented maximally determinately.

When one perceives the colour of a surface, a causal signal emanating from the surface stimulates one's retina. This in turn sets off a cascade of neural firings, which ultimately results in the generation of a conscious perceptual colour representation, a sort of neural or mental symbol. Colour representations can be typed in a variety of ways; I will assume a system of nonintentional types, the *C-types*, which may be biological or functional types.

On a simple-minded version of a familiar *nomic correlational* approach to the metaphysics of content (Fodor, 1990; Field, 1994), for the tokening of a colour representation *r* to have as its content that a surface with colour *C* is present is for there to be a law of nature that under normal conditions, a perceptual encounter with a surface with colour *C* causes a representation of the same C-type as *r* to be tokened. (The restriction of the law to normal conditions allows for misrepresentation, of the sort involved in perceptual illusion or hallucination.) Nothing in this simple-minded approach rules out exact perceptual representation of colour. If there is some maximally determinate colour *C* such that there is a law of nature that under normal conditions, a perceptual encounter with a surface with colour *C* causes a representation of the same C-type as *r* to be tokened, the tokening of *r* is exact: it represents the presence of *C*, a maximally determinate colour, maximally determinately.

But there is never any such law of nature. Like any signal path in the natural world, the signal path in human perception is infected with *noise*: as the signalling process leading from the reflection of the ambient illumination by the colour to the tokening of the colour representation unfolds from event to event, the character of each event in the process is not determined by the events in the process that precede it.[11] Perhaps the unfolding of the process is partly influenced by quantum

[10] At least, I am assuming the first conjunct for simplicity's sake; so long as the properties are represented significantly less determinately than they in fact are, the solution goes through. I take it that indeterminate representation of a property that is determinate and represented as such is coherent: John's height in feet is (perhaps) a real number, and is represented by a certain system as one of the real numbers in a certain range. I won't speculate as to *how* we manage to represent colours as maximally determinate (or as more determinate than the way we represent them); it seems that we do.

[11] I don't know how signalling processes are to be individuated. If any naturalistic theory of content is to be viable, there must be *some* way of individuating them, if only the highly metaphysically committed manner of individuation according to which signalling processes are a natural kind. The upcoming argument for nondetermination of the end of the process by its beginning assumes only that signalling processes are quantum-indeterminate if anything is, and that signalling processes take in significantly less than, and are not isolated from, the entire world.

indeterminacy. Even if not, the process is not isolated from the world outside. These influences external to the process are sufficiently complex that their influence can be treated as a form of randomness for the purposes of understanding the process. Consequently, which C-type is tokened at the conclusion of the process is partly the result of the colour of the surface which initiates the process and partly the result of random noise.[12]

Because of this influence of random noise, a more likely candidate for a law of nature describing colour perception is that under normal conditions, the C-type of a colour representation tokened resulting from a perceptual encounter with an entity with maximally determinate colour C is given by a certain probability density function P'_C with arguments C-types (standardly, P' is viewed as being a so-called 'bell curve'); or, since a 'backward-looking' statement of the law is more useful for the metaphysics of content, that under normal conditions, the maximally determinate colour of an entity a perceptual encounter with which resulted in a colour representation of C-type x is given by a certain probability density function P_x with arguments maximally determinate colours (standardly, also a bell curve).

One approach to modifying the simple-minded theory to square with this law of nature would be to complicate semantic theory by including a probabilistic element.[13] But the present approach to indiscriminability requires a traditional, nonprobabilistic approach to semantics.

Now, an approach on which maximally determinate colours are represented maximally determinately is not *inconsistent* with the probabilistic law of nature. One such approach would take the tokening of a colour representation of C-type x as the result of a process initiated by the colour of a to represent that the colour of a is at the point of the colour solid at which the bell curve P_x reaches its 'peak'. This colour is, in a weak sense, the maximally determinate colour most likely to have been the cause of the tokening (if circumstances are normal).

[12] The classic treatment of perceptual noise is Swets 1964, which seems to predict that we are conscious of noise, as assumptions about the signal/noise ratio apparently influence perceptual content. There are a number of causes of noise, including: quantum indeterminacy influencing how many photons of what energy level are reflected from a surface; random rates of fatigue of the retinal cells; the random sampling rate of opponency channels in the retina; random quivering of the eye; and random patterns of firing in the primary visual cortex. On the last, see Hardin 1988, pp. 171–3; for the rest thanks to James Cutting.

[13] For instance, to introduce a 'semantic density' function, such that the content of representations of C-type x are somehow given by a function P_x with arguments maximally determinate colours.

Still, this approach seems to be intuitively unappealing. I suspect that the reason for this is that we regard *charitable interpretation* as a constraint on a plausible semantic theory.[14] What exactly is required by charitable interpretation is a question well beyond the scope of this paper; still, it seems plausible that any reasonable understanding would have the consequence that given a choice between two semantic theories, one of which assigns a probability of zero to the correctness of a certain type of representation, and one of which assigns a higher probability, the latter is to be preferred, *ceteris paribus*.

Charity pulls away from exactness. If there is no pull toward exactness, charity might pull all the way up to a theory on which perceptual representations represent maximally determinate colours maximally determinably. Still, it is clear that perception *sometimes* enables discrimination of colours; this pulls back down, toward exactness. The pull between charity and discrimination somehow balances out in such a way that maximally determinate colours are represented less than maximally determinately and less than maximally determinably. This is to say that colour representation is inexact.[15] [16] [17]

1.2.2 An argument from phenomenology

The previous argument is directed at those who believe that the possession of a certain colour content by an experience is naturalistically constituted. If colour contents are ultimately the product of probabilities of outcomes of signalling processes of certain types, then, since it is undeniable that these processes are influenced by random noise, it is very

[14] Charity seems to be a prime motivation for the causal correlational approach: if there is a *law* that C's represent *C* under normal circumstances, then incorrect representation will occur only under abnormal circumstances; and abnormality provides an excuse.

[15] Determinacy could be increased without reducing the probability of correctness by 'sampling' or basing the content on the values of a number of representations. Hardin (1988, pp. 174–81) argues that the possibility of sampling shows that the confidence margin can be reduced to zero, so that indiscriminability is transitive after all. But as Williamson (1990, p. 12) notes, Hardin's sampling procedure requires observations of arbitrarily long duration; the notion of perceptual indiscriminability at issue in this paper involves observations of limited duration.

[16] The simple-minded theory is wrong for other reasons, some of which are canvassed in the papers by Fodor and Field. Still, I can't see how complicating the general picture would save maximally determinate representation.

[17] The friend of inexactness may avoid this argument by pressuring its use of charity. Perhaps charitable interpretation requires less than assessment as veridical; rather only to minimize the amount by which experiences are assessed as nonveridical. A view on which noise induces slight illusions would square with this principle. While the approach I sketched is somewhat more in line with traditional understandings of charitable interpretation, assessing these competing approaches is not the business of this paper; hence, unfortunately, I will have to leave discussion of this issue here.

hard to see how to avoid concluding that colours are represented less than maximally determinately.

For those who are less friendly to naturalism, I offer a different case: phenomenologically, it strikes one as if one perceives colours inexactly due to random noise. To careful introspection, there appear to be unavoidable, constant, random, tiny shifts in every aspect of the phenomenal character of every part of a perceptual experience. Gaze fixedly at a white wall of as constant an illumination as possible.[18] Careful introspection will reveal *phenomenal noise*: a subtle 'flickering' or 'crepitation', of the sort manifest in a snowy, poorly tuned rabbit-ears TV. The phenomenal character of your visual experience (however brief) of any area of the wall (however small) is inconstant both across the space of the wall and through the interval of time during which you experience the wall. It would not be correct to characterize this variation of phenomenal character either as involving apparent change in the colour or the illumination of the wall: there is no such apparent change. Rather, the nature of the properties which are apparently varying is not revealed to introspection: they do not clearly strike one either as being instantiated in the external world, or as being in one's own mind, brain, or what have you, or as involving variation in relations between the external world and one's mind.

Phenomenal noise is a sort of variation in the phenomenal character of a visual experience. Does the variation consist in variation in the external properties ostensibly presented in an experience, or in variation in the features of an internal sense-datum 'given' in the experience, or in some further sort of variation in the experience itself, a sort that cannot be reduced to variation of either of these other types? According to Moore (1903), experience is 'transparent' or 'diaphanous'; although Moore's discussion is notoriously obscure, a highly plausible interpretation regards Moore as claiming that any phenomenal character is a property concerning which external features are ostensibly presented in the experience, or which internal features are given in the experience. If Moore is correct here, as I am inclined to think he is, then phenomenal noise is variation of one of the first two sorts. In my view, it is variation of the second sort: a constant, maximally determinate colour is indeter-

[18] Or anything else, for that matter. The white wall facilitates description of the case, since there is no worry about synchronic or diachronic variation in its colour or illumination properties. For things which vary in these ways, phenomenal noise is present in the form of synchronic and diachronic variation not attributable to apparent synchronic or diachronic variation in the colour or illumination of the thing. Phenomenal noise seems to be invariant in its texture density no matter how apparently distant seen things are, whereas texture density of the illumination or surfaces of things increases with apparent distance.

minately 'seen in' a fluctuating sense-datum, in much the way that when watches a snowy TV, the constant, maximally determinate colours of objects presented on the TV are indeterminately seen in the fluctuating pattern of colours on the TV screen.[19] Still, a 'representationalist' view is available on which the variation is of the first sort is available: for instance, a representationalist might take phenomenal noise to be variation in the region of colour space in which an experience represents the colour of a thing as being located.[20] The sense-datum theorist can accept this representationalist conception, explaining the posited variation in terms of variation in internal features; the representationalist would take this sort of variation as not to be explained in terms of more basic aspects of phenomenology. Nothing in this paper hangs on which of these two interpretations of phenomenal noise is correct.

When I say that this phenomenal noise *can* be detected introspectively, I don't mean that everyone attends to it. On the contrary, it is easy to ignore, difficult to attend to. Nor does the fact that people do not always attend to it indicate that it is not a pervasive feature of phenomenal character: we attend to very little, so unless phenomenal character is vastly more impoverished than we pretheoretically take it to be, there is much in phenomenal character to which we do not attend. Nor does the fact that people do not always attend to it mean that it has no influence on their perceptually grounded judgements (including discriminatory judgements). Gibsonian psychologists of perception (cf. Bruce and Green 1985, Pt. III) describe aspects of phenomenal character which influence such judgements (e.g. the direction of 'flow' as being the source of perceptual content concerning motion, or relations to the texture gradient of portions of the 'ground' as being the source of perceptual content concerning distance), but which take great insight to recognize as having this influence.

A psychophysics textbook characterizes noise as 'random variation in a sensory channel' which 'interferes with something you are supposed to' see (Levine and Shefner, 1991, pp. 27, 668). This strikes me as a good characterization of phenomenal noise. When I find myself in a case like Nelson's Predicament, I have reactions like: I would be able to tell *exactly* what the colour of the paint chip is … if this blasted random

[19] See Wollheim 2003 for further discussion of seeing in; I discuss this version of the sense-datum view in my 'Seeing into Sense-Data'.

[20] For more on representationalism, see *inter alia* Harman 1990. Alternatively, the literature stemming from Shoemaker 1994 suggests a general approach to such cases, as involving representation of relations external entities bear to one.

flickering didn't interfere! and: I would be able to tell whether the colour of this chip differs from the colour of that chip … if only it weren't for this obscuring unordered crepitation! Phenomenal noise seems to be random, and seems to interfere with perception of exactly which maximally determinate colours are in the surfaces we see. The frustrating truth about our perceptual condition is revealed by careful introspection: because phenomenal noise interferes, we, like Nelson, can never tell *exactly* what colour paint we are getting.[21]

1.2.3 Regions and phenomenal qualities

A natural supposition at this point would be that phenomenal noise is nothing more than the manifestation in phenomenal character of the noisiness of the perceptual signal path. Sharpening it requires a notion of the *phenomenal quality of the presenting of the colour of* a *in* e, or $Q(a, e)$. This is an aspect of the total phenomenal character of *e*. *e* is structured spatially (and perhaps temporally), and part of the total phenomenal character of the presenting of the colour of *a* in *e* involves that presenting's spatial (and temporal) relations to presentings of other colours. By the phenomenal *quality* of this presenting, I mean to ignore these structural aspects of the phenomenal character of the presenting, and focus exclusively on its *purely qualitative* aspects. Since the C-type flickers and the phenomenal quality flickers, the natural supposition is

[21] Williamson (1994, Sect. 6.4; Sect. 8, esp. 8.3) endorses a more general imprecision in our cognitive capacities. This yields 'margin for error' principles like if one knows that *a* is red, then anything similar in colour to *a* is red (principle (2), p. 181); more generally, if one knows that *x* is *F* then anything similar to *x* in dimensions relevant to status as *F* is also *F*. The argument is that if one can't discriminate between the colour of *a* and the colour of *b*, one's belief that *a* is red can't be reliable unless *b* is red.

Here's how things look on my view. If one knows on the basis of normal, optimal *e* that *a* is red, then $R(a, e) \subseteq$ RED, the region containing all and only the maximally determinate shades of red. This is compatible with there being rare circumstances in which the colour of *a* butts up against the edge of both RED and $R(a, e)$. If this happens, and if it is additionally the case that, for normal, optimal *e*, if $R(a, e) \subseteq$ RED, then one knows on the basis of normal, optimal *e* that *a* is red, Williamson's (2) would go wrong.

Williamson's argument for his margin for error principles assumes that *a*'s and *b*'s colours are indiscriminable *tout court*; and, as I shall argue below in section 2.2, my model fails to predict that there are any instances of perceptual indiscriminability *tout court*.

Later, Williamson argues for margin for error principles in general on the following grounds: 'since a belief constitutes knowledge only if it is reliable enough, the belief that a general condition obtains in a particular case constitutes knowledge only if the condition obtains in all cases similar enough in the relevant respects to achieve the required level of reliability' (pp. 226–7). Processes are reliable in the first instance; a reliable belief is a belief formed on a reliable process. In my case, one's belief that *a* is red is formed on the reliable process of endorsing the content of normal, optimal perception, where inexact contents of normal, optimal perceptual experiences reliably are entailed by the truth. Given this, my view coheres with Williamson's argument, if the relevant respect of similarity is inherited from the noisily reliable tracking of maximally specific colours by perceptual representations.

thus that $Q(a, e)$ determines and is determined by the C-type of the representation tokened in e presenting the colour of a—at least, it is natural to suppose this when e is normal and optimal.

Since, for unaided e, the analogous supposition concerning $R(a, e)$ is forced on us by the simple-minded metaphysics of content, the natural supposition implies that:

(ℸ) If the colour of a is presented in e, and the colour of $a^\star$ is presented in $e^\star$, where both e and $e^\star$ are normal, optimal, unaided experiences; then

$R(a, e) = R(a^\star, e^\star)$ just in case $Q(a, e) = Q(a^\star, e^\star)$.

Since it is *prima facie* plausible and I know of no reason to reject it, I put (ℸ) forth as a conjecture. If it is correct, then just as my analysis allows for indiscriminability without sameness of representation, it allows for indiscriminability without sameness of phenomenology. My experience of a's colour flickers this way; my experience of b's colour flickers that different way; still, the colours are indiscriminable.[22]

1.2.4 The elusiveness of inexactness

The inexactness of perceptual representation of colours is easily ignored, and therefore unlikely to be brought to bear in our philosophical thinking about perceptual indiscriminability. For most purposes, fairly crude perceptual discriminations of colours suffice: such as for selecting an orange from among a bowl of mixed citrus, or locating one's rental car as the taupe one in an aisle of the parking lot occupied by cars of varying colours. The great many gradations we see in those colours due to subtle effects of light and shadow leave the impression that we could make discriminations finer than those we have as yet made. It is natural to project this impression out indefinitely, to suppose that no matter how fine we have gotten, we can always get finer; that is, to suppose that colour perception is exact. It is only in highly rarefied cases, such as those we encounter when wearing our decorator/graphic designer/connoisseur hats, that we butt up against the limits of the projectability of the impression.

Sometimes, when we do this, we notice phenomenal noise. Ordinarily, we don't: it is unimportant to our purposes, so we ignore it. Phe-

[22] Analogous cases for particulars are unexceptional: if one experiences Tweedledee directly and reflected in a red-tinted mirror, one's direct experience of Tweedledee and one's mirror-assisted experience of Tweedledee would differ, and discriminably so, in their phenomenology, though Tweedledee would not be discriminable from himself; alternatively, if it is Tweedledum reflected in the mirror, one might still be unable to discriminate Tweedledee from Tweedledum.

nomenal noise explains the inexactness of colour perception; without an explanation of the falsity of the thesis that colour perception is exact, the natural tendency is to continue to believe it.[23]

1.3 Looking the same

The vernacular 'looking the same' is often invoked in discussions of perceptual indiscriminability: for example, it is said that for *a*'s and *b*'s colours to be indiscriminable is for them to look the same, or for *a* and *b* to look the same in respect of colour.

My apparatus provides four interpretations of the claim that *a*'s and *b*'s colours look the same in *e*. These are: (1) it is incompatible with the content of *e* that *a*'s and *b*'s colour differ—so that $R(a, e)$ and $R(b, e)$ are maximally determinate and identical; (2) that *a*'s colour and *b*'s colour are represented the same way in *e*—so that $R(a, e) = R(b, e)$; (3) that the phenomenal quality under which *a*'s colour is presented in *e* is the same as the phenomenal quality under which *b*'s colour is presented in *e*—so that $Q(a, e) = Q(b, e)$; (4) that *a*'s colour and *b*'s colour are indiscriminable on the basis of *e*—so that $R(a, e)$ overlaps $R(b, e)$.

These interpretations differ significantly in their logical and modal properties. The implication relations among these interpretations are $(1) \Rightarrow (2) \Leftrightarrow (3) \Rightarrow (4)$; $\neg((4) \Rightarrow (3))$; and $\neg((2) \Rightarrow (1))$. The relation involved in (1)–(3) is identity of something or other, a transitive relation; that involved in (4) is overlap, which is intransitive. (1) requires exact representation of colours, which is medically impossible. (2) and (3) are medically possible, but highly unlikely: if there are infinitely many C-types, then the occurrence of either has probability zero; if there are merely a great many C-types, then the occurrence of either is merely unlikely.[24] (4) is medically possible, and occurs in rarefied environments. These differences make it advisable to avoid the 'looks the

[23] A further source of this elusiveness, especially in theoretical contexts, is likely a failure to distinguish sense and reference: *what* is represented is a maximally determinate colour; *how* it is represented is less than maximally determinate.

Does my approach to indiscriminability of colour extend to 'higher-order' indiscriminability — that is, if there is inexactitude in the boundaries of the region *e* represents the colour of *a* as inhabiting, does my approach extend to it? But cases like Nelson's Predicament provide no obvious reason to suppose there is higher-order inexactitude. Since Nelson's experience flickers, there is a clear introspectable difference between the phenomenal characters (with respect to colour) under which the colours of a' and a'' are presented in e_1; hence they are discriminable; hence, by (↯), the regions they are presented as located in differ. Whether a subject can know this on the basis of the experience alone depends on such things as the subject's ability to make judgements concerning the region in which a certain colour is represented as being; plausibly this ability is not required for, and is rarer than, the ability to make judgements concerning the colours of things. Still, a subject's lack of this ability is no fault of the experience, and it seems extremely plausible that a subject in possession of the ability would be able to discriminate the relevant regions.

same' locution in philosophical discussions of perceptual indiscriminability in which these logical or modal properties might be at issue.

If ordinary language does not distinguish these interpretations, this is because, as Strawson (1976, pp. 24–5) has pointed out, 'the same' can apply to two things when they are different, just so long as they are are sufficiently similar for the purposes of the conversants; accordingly, 'looks the same' can apply to two things which look different, just so long as the ways they look are sufficiently similar for one's purpose. Most uses of 'looks the same' do not occur in contexts in which conversants have available special instrumentation that makes discriminations finer than those made by the eye; for the conversants' purposes, perceptual indiscriminability would be similarity enough.

2. Saving exact representation

In this section, I will examine the relative merits of my line on Nelson's Predicament and a competing package originating in a paper by Jackson and Pinkerton (1973). The first part of the package is a positive view: colour representation is exact; in Nelson's Predicament, for at least one of the colours, it is represented in one experience as being at one point of the colour solid and in another experience illusorily as being at a distinct point; owing to exactitude, the indiscriminability-making relation ~ collapses into identity and is therefore, like perceptual indiscriminability *tout court*, transitive. The second part of the package consists of objections to nontransitivity of ~; the third of a strategy for undermining the apparent bearing of Nelson's Predicament on nontransitivity of ~. I have already argued against the first part of the package (by arguing for the denial of the exactness thesis fundamental to it); I will now argue against the second and third parts.

[24] This affords a reply against the following sorites argument: for each pair, the looks (in senses (2) or (3)) of the members of the pair are the same; by transitivity of 'the same', the looks of the first and last members of the series are the same; but the looks of the first and last members are different. A common reply due to Jackson and Pinkerton (1973; see also the authors discussed in fn. 27, below) is to deny validity on the grounds that there is no assurance that the looks of the members of the pairs across experiences have not altered. This will inevitably seem somewhat speculative, and leave one wondering whether such shiftiness is *required* for the relevant situation to arise (cf. Mills 2002, which I discuss in section 3.3, below).

On my treatment, the probability is zero that for any randomly selected experience, the looks (in senses (2) and (3)) of the members of a pair of colours will be the same; this fact is manifest upon reflection, given the apparently random character of phenomenal noise. Hence we have no reason to believe that the sorites argument is sound because we have no reason to believe it has true premisses, not just because we have no reason to believe it is valid.

2.1 Objections to nontransitivity

2.1.1 Jackson and Pinkerton

Jackson and Pinkerton (1973, p. 271) purport to derive an absurdity from the nontransitivity of ∼:

> [T]he suggestion that *A* might look to be the same colour as *B*, *B* might look to be the same colour as *C*, while *A* looks to be a different colour from *C*, *to one and the same person at one and the same time*, is inconsistent. As *A* and *C* *ex hypothesi* look to be different colours, looking to be the same colour as *A* will be distinct from looking to be the same colour as *C*; therefore, the suggestion involves one object *B*, looking to have two different colours at the same time to the same person, which is impossible.[25]

I reconstruct this line of reasoning as follows.

(i) Suppose a case like Nelson's Predicament shows that the colours of *A* and *B* are indiscriminable, and the colours of *B* and *C* are indiscriminable, but the colours of *A* and *C* are discriminable.

(ii) So the colours of *A* and *B* look the same, and the colours of *B* and *C* look the same, but the colours of *A* and *C* look different.

(iii) So, considering the colours things look to have: those for *A* and *B* are the same, and those for *B* and *C* are the same, but those for *A* and *C* are different.

(iv) But this implies that there are two colours *B* looks to have, which is impossible.

Discussion in section 1.3, above, revealed four different readings of the claim that the colours of *a* and *b* 'look the same' in *e*. On readings (1)–(3), the argument does not get to stage (ii), since on my analysis, indiscriminability is insufficient for looking the same on any of these readings. On disambiguation (4), stages (ii) and (iii) run as follows:

(ii-4) So, for any normal, optimal, unaided *e*, *R*(*A, e*) and *R*(*B, e*) overlap, as do *R*(*B, e*) and *R*(*C, e*), but *R*(*A, e*) and *R*(*C, e*) exclude one another.

(iii-4) So, considering the region the colour of *A* might be in (consistent with the content of a normal, optimal, unaided experience), and the region the colour of *B* might be in (consistent with the content of a normal, optimal, unaided experience), and the re-

[25] Jackson (1977, p. 114) and Graff (2001, pp. 914–6) press the same argument.

gion the colour of *C* might be in (consistent with the content of a normal, optimal, unaided experience):

the first and the second overlap, and the second and the third overlap, but the first and the third don't overlap.

But if colour perception can be inexact, (iii-4) provides no support for (iv).

2.1.2 Graff

Let a *phenomenal sorites series* be a series of objects numbered 1 to *n* such that object *i* looks the same in colour as object *i* + 1, but such that, for some colour *C*, object 1 but not object *n* looks *C*; let a *same-appearance claim* be a claim that if *a* and *b* look the same in colour, then if one looks *C* so does the other. Graff (2001, Sect. 1) argues against nontransitivity of ~ as follows:

(v) if *looking the same in colour* is nontransitive, there could be a phenomenal sorites series;

(vi) if there could be a phenomenal sorites series, some same-appearance claim is not true;

(vii) same-appearance claims are 'truisms' (p. 907), such that one who rejected one 'would not merely seem to be plainly mistaken, but also to be in a state of confusion' (p. 909).

Once again, the argument is irrelevant to indiscriminability on readings (1)–(3) of 'look the same'.

Suppose the argument is read as per (4): then, in order to assess the premisses of the argument, we need an interpretation of the notion of a same-appearance claim along the lines of (4), which in turn requires an interpretation of the notion of *a*'s looking *C*. Compare '*a* looks red': this seems to mean that the colour of *a* is represented as being a maximal determinate of red. Generalizing, for *a* to look *C* (in *e*) is for $R(a, e) = C$.

Then the argument runs:

(v-4) if indiscriminability is nontransitive, there could be a phenomenal sorites series;

(vi-4) if there could be a phenomenal sorites series, for some *a*, *b*, *e*, the colour of *a* and the colour of *b* are indiscriminable on the basis of *e*, though $R(a, e) \neq R(b, e)$.

(vii-4) if two colours are indiscriminable on the basis of *e*, for any region of colour space, one is represented by *e* as being in it if the other is: in effect, if the colours of *a* and *b* are indiscriminable on the basis of *e*, $R(a, e) = R(b, e)$.

That is, (vii-4) takes indiscriminability to require sameness of representation, or 'looking the same colour' on reading (2). But a central principle of my treatment of indiscriminability is that it does not require sameness of representation; similarity of representation sufficient for overlap is enough. Only on the assumption of exactness does overlap require identity. In plain language, (vii-4) is: if the colours of two things are indiscriminable, one of them looks red if the other does (and so forth). Even in plain language, this isn't plausible: what is plausible is that if the colours of two things are indiscriminable, if one looks red the other doesn't look not red. An experience noncommittal as to whether the other is red won't always allow discriminating its colour from a red colour.[26]

2.2 Undermining the case for nontransitivity

The undermining strategy purports to establish that the data from Nelson's Predicament do not demonstrate that 'perceptual indiscriminability is nontransitive'. The data from Nelson's Predicament are that the colours of a' and a'' are indiscriminable on the basis of e_1, that the colours of a'' and a''' are indiscriminable on the basis of e_2, and that the colours of a' and a''' are discriminable on the basis of e_3. As Graff puts the point, to conclude nontransitivity on these data 'would be like concluding that 'is taller than' is not transitive from the following: Tim is taller than Michael (in 1980); Michael is taller than David (in 1990); and Tim is not now taller than David' (p. 913).[27] I located two readings of this claim in section 1.1, above: the weaker claim that the indiscriminability-making relation $\sim$ is nontransitive, and the stronger claim that perceptual indiscriminability *tout court* is nontransitive.

[26] While (vii) is implausible on reading (4), (v) and (vi) are plausible on any reading. Concerning (v): if *looking the same in colour as* is nontransitive, then there could be a series of objects numbered 1 to *n* such that object *i* looks the same in colour as object $i + 1$, but object 1 does not look the same in colour as object *n*. In order to establish that such a series is a phenomenal sorites series, the principle is required that if *a* does not look the same in colour as *b*, there is some colour *C* such that *a* looks *C* but it is not the case that *b* looks *C*. This seems indisputable.

Concerning (vi): that object *n* in the series looks *C* (*contra* the stipulation) follows from the assumption that object 1 looks *C*, the assumption that for each *i*, object *i* looks the same in colour as object $i+1$, and the same-appearance claim that if *a* looks the same in colour as *b* then if one looks *C* so does the other.

[27] The observation originates with Jackson and Pinkerton (1973, p. 270); cf. Jackson 1977, p. 113; Williamson 1994, p. 174; Raffman 2000, pp. 163–4, 169–71; Graff 2001, pp. 913, 934.

Against the former, the significance of the Tim–Michael–David case seems to be that whether *a* is taller than *b* at a time depends on the nature of the heights *a* and *b* have at the time—specifically, whether the former bears > to the latter—and the case doesn't show that > is nontransitive, since for this, facts are required concerning how Tim's heights in 1980 and now are related, how Michael's heights in 1980 and 1990 are related, and how David's heights in 1990 and now are related.

Analogously, whether the colours of *a* and *b* are indiscriminable on the basis of *e* depends on which regions $R(a, e)$ and $R(b, e)$ are: whether the former bears some relation ~ to the latter. Nelson's Predicament doesn't show that ~ is nontransitive—that the indiscriminability-making relation is nontransitive. Further claims are required to show this: $R(a', e_1) \sim R(a', e_3)$, $R(a'', e_1) \sim R(a'', e_2)$, and $R(a''', e_2) \sim R(a''', e_3)$. If ~ is identity, or if the *R* are of radius zero, these further claims can't (by math) all be true; if ~ is overlap, and the *R* are of positive radius, they can (so far as math is concerned). Whether it is *medically* possible for them all to be true depends on how many C-types there are; a continuum of C-types would suffice, but all that is necessary are enough C-types relative to the radius of the *R*; a reasonably permissive (though not implausibly so) recombination principle for experiences is also needed.

Against the latter, stronger reading of the claim that perceptual indiscriminability is nontransitive, namely that perceptual indiscriminability *tout court* is nontransitive, the worry seems to be this: indiscriminability *tout court* is indiscriminability on the basis of any (possible, normal, optimal, unaided) experience. Of course Nelson's Predicament by itself shows nothing about indiscriminability *tout court*: for this, I needed the assumption that what goes for the discriminability status of two colours on the basis of any such experience goes for their discriminability status on the basis of all such experiences. And perhaps this assumption, like the assumption that what goes for David and Tim's relative tallness at some time goes at all times, is unwarranted.

My explanation of inexactness suggests a case against the assumption. Due to noise, for any maximally determinate colour *C* and any C-type *T*, there is some possible experience in which *C* is perceived via the tokening of a representation which is *T*. Only some such experiences will count as veridical: for this, it is necessary and sufficient that the region R_T of radius δ of colour space associated with representations of type *T* includes *C*.[28]

Still, suppose that the maximally determinate colours of *a* and *b* are distinct. Then, if δ is of a reasonably small value, there are regions R_a and R_b of radius δ such that the former includes the the colour of *a* but not the colour of *b* and the latter includes the colour of *b* but not the colour of *a*. If there are enough C-types, for some possible experience *e* of the colour of *a* and the colour of *b*, $R(a, e) = R_a$ and $R(b, e) = R_b$. *e* is veridical, since the colour of *a* is in $R(a, e)$ and the colour of *b* is in $R(b, e)$; there is no other reason to doubt the optimality or normality of *e*, and *e* is unaided. Moreover, *e* is correct only if the colour of *a* has a property the colour of *b* lacks. This might seem to suffice for *e*'s warranting a judgement of the distinctness of the colours of *a* and *b*; if so, these colours are discriminable on the basis of *e*. The only assumption about these colours is that they are distinct; the argument would seem to show that for any distinct pair of colours, they are not perceptually indiscriminable *tout court*.

Two reactions to this argument strike me as viable. The first is to note that one of its premisses is a certain non-trivial *sufficient* condition on knowledge (that, in optimal, normal conditions, judging that *p* on the basis of a perceptual experience incorrect unless *p* suffices for recognizing, and thus knowing, that *p*). This premiss is highly debatable, since non-trivial sufficient conditions on knowledge are difficult to adjudicate. By contrast, trivial sufficient conditions on knowledge, such as that seeing that *p* or remembering that *p* suffices for knowing that *p*, are a dime a dozen; as are non-trivial necessary conditions, such as its being the case that *p*. But in order to establish the premiss, something like a solution to the Gettier problem might be required.

The second reaction to the argument is to retreat to a less demanding notion of indiscriminability than my indiscriminability *tout court*: for instance, the colour of *a* and the colour of *b* are *perceptually p-indiscriminable* $=_{df}$ the probability that the colour of *a* and the colour of *b* are indiscriminable on the basis of *e* (given that *e* is possible, normal, optimal, and unaided) is greater than *p*. Two colours which are not so similar will be 0-indiscriminable; as their degree of similarity increases toward identity, the value of *p* for which they are *p*-indiscriminable increases toward 1. Whether the notion of *p*-indiscriminability, indiscriminability *tout court*, or just ~, is appropriate to the treatment of this

[28] Raffman (2000) and Graff (2001, esp. p. 926) both present a budget of interesting psychological data which indicates the great shiftiness in our perceptual systems: for example, putting a square next to a lighter square makes the first look darker than it looks when put next to a darker square. A possible way around these data (which I am uncertain whether to endorse) is to dispute their relevance to the argument on the grounds that certain of the situations Raffman and Graff describe induce illusions about colour, so that the experiences are not *optimal*.

or that philosophical problem will need to be assessed on a case-by-case basis.

3. Consequences of nontransitivity

In this concluding section, I urge caution when drawing philosophical consequences out of Nelson's Predicament.

3.1 Armstrong against sense-data

Armstrong (1968, p. 218) appeals to the nontransitivity of ~ in arguing against sense-data:

> 'Exact similarity in a particular respect' is necessarily a transitive relation. Now suppose that we have three samples of cloth, *A*, *B*, and *C*, which are exactly alike except that they differ very slightly in colour. Suppose further, however, that *A* and *B* are *perceptually* completely indistinguishable in respect of colour [in $e^{\star}$], and *B* and *C* are perceptually completely indistinguishable in respect of colour [in $e^{\star\star}$]. Suppose, however, that *A* and *C* can be perceptually distinguished from one another in this respect [in $e^{\star\star\star}$].
>
> Now consider the situation if we hold a 'sensory item' view of perception. [...] [I]t will seem to follow that the two sensory items A_1 and B_1 that we have when we look at the two pieces [*A* and *B*] *actually are identical in colour*. For the sensory items are what are supposed to make a perception the perception it is, and here, by hypothesis, the *perceptions* are identical. In the same way B_1 and C_1 will be sensory items that are identical in colour. Yet, by hypothesis, sensory items A_1 and C_1 are not identical in colour.

Suppose: $R(A, e^{\star}) \sim R(B, e^{\star})$; $R(B, e^{\star\star}) \sim R(C, e^{\star\star})$; $\neg(R(A, e^{\star\star\star}) \sim R(C, e^{\star\star\star}))$; also that $R(A, e^{\star}) \sim R(A, e^{\star\star\star})$; $R(B, e^{\star}) \sim R(B, e^{\star\star})$; $R(C, e^{\star\star}) \sim R(C, e^{\star\star\star})$. Allegedly, the sense-datum theorist individuates 'colours' (better, *colour's*, on the familiar Peacockean terminology for the alleged colour-like properties of sense-data) of sense-data via a principle according to which if $R(a, e) \sim R(b, e)$, *e* consists in part of acquaintance with a sense-datum of *a* and a sense-datum of *b* which are the same in colour′. Armstrong's conclusion surely follows.

But it is more in the spirit of sense-datum theory to begin with a principle for individuating colour′s on which the consequent follows on $Q(a, e) = Q(b, e)$. This implies Armstrong's principle only on the assumptions of (♮) and exactness; I have of course argued extensively against the latter.

3.2 Wright against colour predicates

Wright (1975, pp. 338–9) has appealed to Nelson's Predicament in arguing that colour predicates are inconsistent:

> Colour predicates [...] are in the following sense *purely observational*: if one can tell at all what colour something is, one can tell just by looking at it. The look of an object decides its colour. [...] Since colour predicates are observational, any pair of objects indistinguishable in point of colour must satisfy the condition that any basic colour predicate applicable to either is applicable to both. It is, however, familiar that we may construct a series of suitable, homogeneously coloured patches, in such a way as to give the impression of a smooth transition from red to orange, where each patch is *indiscriminable* in colour from those immediately next to it; it is the non-transitivity of indiscriminability which generates this possibility. So, since precise matching is to be sufficient for sameness of colour, we can force the application of 'red' to all the patches in the series, some of which are not red but orange.

Allegedly, (viii) colour predicates are observational, and (ix) observational predicates are inconsistent. Supposing that colour predicates map onto regions of the colour solid, (viii) seems to have as a consequence the principle (vii-4) from section 2.1, above; the argument to (ix) goes from (v-4), (vi-4), and the nontransitivity of $\sim$. This is paradoxical on the assumption that no significant fragment of natural language is inconsistent.[29]

Unfortunately, (vii-4) is incorrect.[30] This ramifies back to the claim that 'the look of a thing decides its colour': apparently, that for any colour predicate, and any object, one who grasps the predicate can always know just by looking (under normal, optimal conditions) whether the predicate applies to the object. When I apply a colour predicate to a colour, I do so by looking at the colour to determine which colour it is, and apply a predicate appropriate to its nature as I determined it. (viii) requires that I can always tell, just by looking, exactly which colour a colour is. But phenomenal noise interferes with my ability to do this. Perception is manifestly inexact. Why would we adopt a linguistic practice at odds with this?[31]

3.3 Mills against privileged access

According to Mills (2002), one cannot tell when two things look the same to one. As I interpret Mills, he argues as follows:

[29] Dummett (1975, p. 264) defends (ix). It is unclear that this by itself yields paradox. So what if observational predicates are inconsistent? So is 'tonk'. Fortunately, English doesn't contain it.

[30] Would it be bad enough that a central fragment of the language presupposes a falsehood?

[31] This dialectic is taken further in Williamson 1994, Sect. 6.4.

> Assume that Nelson′ is in a predicament which appears introspectively to him to be Nelson's Predicament (so that apparently ~ is nontransitive). Assume that Nelson′'s experience is constant; and that ~ is identity. But then a contradiction follows by the logic of identity. So either (a) the *actual* patterns of discriminability and indiscriminability manifest in Nelson′'s predicament are not as they appear introspectively to Nelson′, or (b) Nelson′'s experience is shifty. But endorsing (b) is not generally adequate, since (Mills strongly intuits) there could be a being with non-shifty experiences who fell into a situation like Nelson′'s predicament. Hence (a); then one can't always tell introspectively whether two colours are perceptually indiscriminable to one. The interior world is not transparent to itself; a radical externalism threatens.

The principle that ~ is identity is plausible only on the assumption of exactness, so a third way out of the contradiction is to reject this. I expect Mills to rejoin that inexactness requires noisiness, which is like shiftiness, and that it is strongly intuited that there could be a being whose perception is *clean*—noiseless—who could fall into a predicament like Nelson′'s, and then go on as before.

But I do not intuit this. On my account, noise is a central source of the nontransitivity of perceptual indiscriminability even under optimal, normal circumstances. Noise blurs subtle differences; for sufficiently similar colours, this yields uncertainty whether they are distinct. If all signals were clean, perhaps only identical colours would be perceptually indiscriminable.[32]

The Sage School of Philosophy BENJ HELLIE
218 Goldwin Smith Hall
Cornell University
Ithaca, NY 14853
USA
benj.hellie@cornell.edu

References

Armstrong, David Malet, 1968: *A Materialist Theory of the Mind*. London: Routledge and Kegan Paul.

[32] Comments by the editor of *Mind* on previous drafts improved the paper immeasurably. Thanks also to Tamar Szabó Gendler and Jessica Wilson for detailed comments and discussion; to David Jehle and audiences at Ohio State University and the University of Toronto; and, for inspiring discussion of her paper on this material, to Delia Graff.

Bruce, Vicki and Patrick Green, 1985: *Visual Perception: Physiology, Psychology, and Ecology*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Dummett, Michael A. E., 1975: 'Wang's Paradox', in his *Truth and Other Enigmas*. Cambridge, MA: Harvard University Press 1978, pp. 248–68. Originally published in *Synthese*, 30, pp. 201–32.

Evans, Gareth, 1982: *The Varieties of Reference*. Oxford: Oxford University Press.

Field, Hartry, 1994: 'Deflationist Views of Meaning and Content', in his *Truth and the Absence of Fact*. Oxford: Oxford University Press 2001, pp. 104–40. Originally published in *Mind*, 103, pp. 249–85.

Fodor, Jerry A., 1990: *A Theory of Content and Other Essays*. Cambridge, MA: The MIT Press.

Goodman, Nelson, 1951: *The Structure of Appearance*, 1st edition. Cambridge, MA: Harvard University Press.

Graff, Delia, 2001: 'Phenomenal Continua and the Sorites'. *Mind*, 110, pp. 905–35.

Hardin, Clyde L., 1988: *Color for Philosophers: Unweaving the Rainbow*. Philadelphia: Hackett.

Harman, Gilbert, 1990: 'The Intrinsic Quality of Experience', in his *Reasoning, Meaning, and Mind*. Oxford: Oxford University Press 1999, pp. 244–61. Originally published in *Philosophical Perspectives*, 4, pp. 31–52.

Hecht, Heiko, Robert Schwartz and Margaret Atherton (eds), 2003: *Looking into Pictures*. Cambridge, MA: The MIT Press.

Jackson, Frank, 1977: *Perception: A Representative Theory*. Cambridge: Cambridge University Press.

Jackson, Frank and R. J. Pinkerton, 1973: 'On an Argument Against Sensory Items'. *Mind*, 82, pp. 269–72.

Levine, Michael W. and Jeremy M. Shefner, 1991: *Fundamentals of Sensation and Perception*, 2nd edition. Pacific Grove, CA: Brooks/Cole Publishing Company.

Lewis, David, 1996: 'Elusive Knowledge', in his *Papers in Metaphysics and Epistemology*, pp. 248–68. Originally published in *Australasian Journal of Philosophy*, 74, pp. 418–45.

Lewis, H. (ed.), 1976: *Contemporary British Philosophy*, 4th series. London: George Allen and Unwin.

Mills, Eugene, 2002: 'Fallibility and the Phenomenal Sorites'. *Noûs*, 36, pp. 384–407.

Moore, George Edward, 1903: 'The Refutation of Idealism'. *Mind*, 12, pp. 433–53.

Pryor, James, 2000: 'The Skeptic and the Dogmatist'. *Noûs*, 34, pp. 517–49.
Raffman, Diana, 2000: 'Is Perceptual Indiscriminability Nontransitive?' *Philosophical Topics*, 28, pp. 153–75.
Shoemaker, Sydney, 1994: 'Phenomenal Character'. *Noûs*, 28, pp. 21–38.
Strawson, Peter F., 1976: 'Entity and Identity', in his *Entity and Identity*, pp. 21–51. Oxford: Oxford University Press 1997. Originally published in Lewis 1976.
Swets, John A., 1964: *Signal Detection and Recognition in Human Observers: Contemporary Readings*. New York: John Wiley and Sons.
Williamson, Timothy, 1990: *Identity and Discrimination*. Oxford: Blackwell.
——1994: *Vagueness*. London: Routledge.
——1995: 'Is Knowing a State of Mind?' *Mind*, 104, pp. 533–65.
Wollheim, Richard, 2003: 'In Defense of Seeing In'. In Hecht, Schwartz, and Atherton 2003, pp. 3–16.
Wright, Crispin, 1975: 'On the Coherence of Vague Predicates'. *Synthese*, 30, pp. 325–65.