Bressan, P. (2007). Dungeons, Gratings, and Black Rooms: a Defense of Double-Anchoring Theory and a Reply to Howe et al. (2007). *Psychological Review, 114,* 4, 1111-1115.

Dungeons, Gratings, and Black Rooms:

a Defense of Double-Anchoring Theory and a Reply to Howe et al. (2007)

Paola Bressan
Università di Padova

The double-anchoring theory of lightness (P. Bressan, 2006b) assumes that any given region belongs to a set of frameworks, created by Gestalt grouping principles, and receives a provisional lightness within each of them; the region's final lightness is a weighted average of all these values. In their critique, P. D. L. Howe, H. Sagreiya, D. L. Curtis, C. Zheng, and M. S. Livingstone (2007) (a) show that the target's lightness in the dungeon illusion (P. Bressan, 2001) and in White's effect is not primarily determined by the region with which the target is perceived to group and (b) claim that this is a challenge to the theory. The author argues that Howe et al. misinterpret grouping for lightness by equating it with grouping for object formation and by ignoring that lightness is determined by frameworks' weights and not by what *appears* to group with what. The author shows that Howe et al.'s empirical findings, together with those on grating induction and all-black rooms that they cite as problematic, actually corroborate, rather than falsify, the double-anchoring theory.

Keywords: dungeon illusion, grating induction, White's effect, reverse contrast, anchoring theories

In their critique of my double-anchoring theory (Bressan, 2006b) Howe, Sagreiya, Curtis, Zheng, and Livingstone (2007) raise four assorted points. First, in the dungeon and White's illusions, the target's lightness is not primarily determined by the region with which the target is *perceived* to group. Second, in the dungeon illusion, a 4-element context affects the lightness of the target less than a 56-element context. Third, "the lightness of a point is most affected by the luminance of nearby points". Fourth, the all-black miniature room in Gilchrist and Jacobsen's (1984) experiment was matched to a gray, rather than a white, on a Munsell chart. Howe et al. believe that these four points represent problems for the theory.

Does DAT claim that the target's lightness is primarily determined by the region with which the target is *perceived* to group?

In DAT, a target typically belongs to multiple frameworks; Howe et al. contend that the strongest (or "dominant") one must correspond to the perceptual group. However, the strength of a framework is defined "as a function, first, of its relative size, articulation, and absolute luminance; second, of the number and type of spatial and photometric grouping factors that make the target belong to it" (Bressan, 2006b, p. 531), which is not the way a perceptual group is defined.

There are many instances in my original article where the dominant framework and the perceptual group do not coincide; one such example is the 5-square Gelb effect (Bressan, 2006b, Appendix B), in which each square in a row of five is perceived to group with the other squares regardless of whether the spotlight illuminating the squares is on or off. However, relative to the peripheral framework (that comprises the target square and the laboratory), the local framework (that comprises the target square and the other squares) is dominant when the spotlight is on (weight ratio 1:0.35), and nondominant when the spotlight is off (weight ratio 1:1.75). This means that when the spotlight is off (i.e., it is replaced by normal room illumination) each Gelb square groups with the room more than with the other squares! This kind of grouping, surely, cannot be mistaken for "which regions are *perceived* to group together".

The concept that frameworks are not the same as the perceptual groups (already expressed in Bressan, 2001) follows directly from their being defined as regions of common illumination, not as objects or groups thereof. In the two domains, the grouping principles have accordingly different relative weights. In the case of object formation (perceptual grouping), the weightier principles are those that represent a better *proxy for object unity*. Because regions that move coherently together are likely to belong to the same object, here the principle of common fate is a powerful one—as shown by its value in creating or breaking camouflage (Metzger, 1936/2006). In the case of lightness assessment, however, the weightier principles are those that represent a better *proxy for shared illumination*. Because regions that move coherently together are not especially likely to be illuminated the same way, here common fate is not very informative (it was not even listed among grouping factors in Bressan, 2006b). Indeed, the biologically most interesting regions that move together are perhaps the different parts of an animal; and these are typically illuminated in a different manner—as shown by the evolution of countershading (lighter pigmentation of ventral regions) in mammals and birds to counterbalance the depth-revealing effects of being lit from above (Thayer, 1896).

A poor grouping factor such as common fate can still tilt the balance if the two conflicting frameworks have essentially equal strengths (as in Agostini & Proffitt, 1993). If strengths are different to start with, as in the dungeon illusion, its introduction may still leave a slight trace, and interestingly the data that Howe et al. present in Figure 3 suggest just that: For each subject, the target is perceived as slightly lighter in Display A than in Display B, and this is exactly what one would predict if grouping by common fate would have some small effect, because in A the target groups better with the black background, and in B with the white grid. Moreover, for each subject except subject DL, the target is perceived as slightly darker in Display C than in Display D, and this is again exactly what one would expect, because in C the target groups better with the white background, and in D with the black grid.

It follows that Howe et al.'s data, rather than representing a challenge to DAT, go, if anything, in the direction predicted by it. Howe et al.'s experiments are, in essence, a low-power replication of a previously reported weak effect (Agostini & Proffitt, 1993). It is odd that these data should be portrayed as evidence against anchoring theories, the only theories that can actually accommodate them.

Can DAT explain the reduced-context dungeon illusion?

Howe et al. open the corresponding section of their critique with a misunderstanding and close it with another. When they say that “a grouping principle cannot be categorized as hard or soft before the context is known”, they are confusing the dimension hard/soft with the dimension strong/weak. The “soft” principles are those that refer to the state of the observer, such as attention or past experience (Bressan, 2006b, p. 530; Bressan, 2006a); the “hard” principles are those that refer to the stimulus, such as proximity, common region, alignment, depth and shape similarity, luminance polarity and similarity—and this distinction holds before and after the context is known.

Howe et al. proceed to claim that, to predict the difference between the full- and reduced-context dungeon illusions, DAT must concede that the same principle can be weak in one context and strong in another. But this is not the case either. In the dungeon illusion, the disks-framework supports reverse contrast and the background-framework supports simultaneous contrast. The direction of the final effect depends on the relative strengths of the two frameworks, and therefore on their relative size, articulation, and absolute luminance. In the reduced-context dungeon, the framework responsible for reverse contrast has been weakened (in both relative size, less area covered by contextual disks, and articulation, 4 vs. 56 contextual disks; because of the larger distance between targets and contextual disks, the grouping principle of proximity is reduced also), whereas the framework responsible for simultaneous contrast has been strengthened (in relative size, more area covered by background). Is the perceptual difference between the two displays a challenge to DAT? I would say, more like an unintended confirmation.

Can DAT explain grating induction?

Contrary to Howe et al.’s assumption, in DAT parts of a scene that have the same luminance but *different surrounds* cannot be considered “identical” (and therefore “treated as a unity”), because their local frameworks are different.

In DAT, grating induction is essentially a variant of White’s effect in which the gray target regions laying on the white and black bars of the grating are joined together (see

Figure 1). DAT explains White's effect (not covered in Bressan, 2006b) by showing that each target patch gives rise to two main frameworks, consisting of the target and either the collinear bar or the flanking bars. In the collinear-framework, the target on the white bar is gray at both the highest-luminance and surround steps; the target on the black bar is, respectively, white and superwhite. Thus, in the collinear-framework, the target on white is darker than the target on black. In the flanking-framework, the target on white (flanked by black bars) is white and superwhite at the highest-luminance and surround steps, whereas the target on black (flanked by white bars) is gray at both steps. Thus, in the flanking-framework, the target on white is lighter than the target on black.

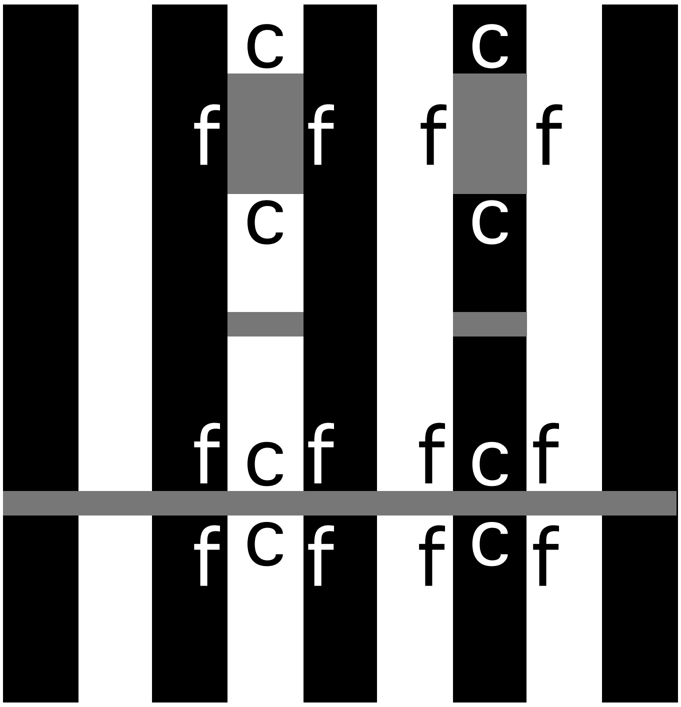

Figure 1. White's effect (top) and grating induction (bottom). In both cases, the gray regions sitting on white look darker than those sitting on black. Each target patch belongs to a collinear-bar framework, *c*, and a flanking-bar framework, *f*. DAT explains either illusion as a combination of anchoring in the collinear-bar framework and anchoring in the flanking-bar framework; because of their opposite luminance polarities relative to the target, these frameworks have opposite effects on the target's lightness. The direction of the illusion depends on the relative weights of the two frameworks; here, the collinear framework is stronger in both cases (see text for details). The two middle patches illustrate a White's effect where the target has the same width as in the top version, but a reduced height.

The flanking-framework is founded on the grouping principle of proximity (in its strong form of adjacency), whereas the collinear-framework is sustained by both proximity (in its strong form of adjacency) *and* good continuation (based, here, on the powerful combination of T-junctions and orientation similarity). The latter framework, that supports simultaneous contrast with the collinear bar, clearly weighs more than the former, that supports simultaneous contrast with the flanking bars; the perceptual outcome is White's effect.

If the target regions of White's display are aligned so as to form a long single stripe (Figure 1, bottom) a grating-induction display arises. The flanking-framework is still founded on the principle of proximity, now in a weaker form—because adjacency is gone everywhere, except at the (virtual) corner regions defined by the newly created T-junctions. The collinear-framework is still sustained by proximity in its strong form of adjacency, but good continuation has been dismantled entirely. As in White's effect the latter framework is clearly stronger than the former, resulting in simultaneous contrast with the collinear bar, that is grating induction. However, the absence of good continuation makes this illusion different from White's effect in predictable ways.

For example, although articulation and absolute luminance do not consistently favor one framework over the other, the weighting factor of relative size can pull in either direction depending on the aspect ratio of the target patch. In DAT, larger frameworks weigh more than smaller ones. With a fixed-height target, a decrease in bar (and thus in target) width will hence weaken the collinear-framework relative to the flanking one, *reducing* the illusion. This has been documented for both White's effect (Kingdom & Moulden, 1991) and grating induction (e.g. Foley & McCourt, 1985).

With a fixed-width target patch, a decrease in target height is expected to weaken the flanking-framework relative to the collinear one, *increasing* the illusion. For grating induction, this prediction is well supported (e.g. Foley & McCourt, 1985). However, in White's effect a decrease in target height (Figure 1, middle) will concurrently weaken the grouping factor of good continuation, by affecting orientation similarity (that depends crucially on the sign, and less crucially on the magnitude, of the horizontal:vertical ratio). Therefore, the strengthening of the collinear-framework due to the weighting factor of relative size will be partly, fully, or even over- compensated by a corresponding abatement due to poorer grouping. This instability is reflected by the discordant results obtained in different laboratories (Kingdom & Moulden, 1991, vs. B. Spehar, personal communication, February 2007) and by the existence of obvious individual differences (e.g., subject FK vs. subject BM: Kingdom & Moulden, 1991, Figure 6). Because good continuation supports collinear grouping in White's effect but not in grating induction, DAT further predicts that the latter will subside sooner than the former when the target region grows taller. This has been empirically observed as well (Zaidi, 1989).

DAT's interpretation of grating induction and White's effect also accounts for hitherto unexplained aspects of these illusions, such as the fact that both disappear when the luminance of the target is either higher than the luminance of the bright bar or lower than the luminance of the dim bar (Spehar, Gilchrist, & Arend, 1995). In these cases, one of the two collinear-frameworks is drastically weakened by the fact that the factors of luminance polarity and similarity, which were neutral in the regular display, now promote grouping with the flanking bars. For example, if the target is the lowest luminance in the display, the patch on the dim bar still groups primarily with the dim collinear bar (like in the regular effect) but the patch on the bright bar now groups primarily with the dim flanking bars (unlike in the regular effect). Within their dominant frameworks (the collinear one in the first case, the flanking one in the second), the targets thus receive the *same* lightness assignment. The relative weights of the two frameworks are accordingly crucial for understanding whether the regular effect will simply diminish, disappear, or reverse. If the framework comprising the dim flanking bars is much larger, as is the case with the typically elongated targets of White's effect, the target on the bright bar will appear lighter than the target on the dim bar; that is, White's effect will reverse. DAT predicts that the larger the ratio between the two sides of the target (i.e., the more elongated the target), the higher the probability that the double-decrement or double-increment White's effect reverses (or, the larger the reverse-contrast illusion).

Illusions based on conflicting frameworks—where the target has an intermediate luminance and photometric grouping factors are hence neutral—are critically sensitive to the overturning of the luminance hierarchy, because the combination of luminance polarity and similarity can prevail over spatial grouping factors (and make even T-junctions ineffective: Bressan, 2001; Bressan, 2006b). In White's display, the result is the inverted-White illusion (e.g. Ripamonti & Gerbino, 2001); in grating induction, the result is "visual phantoms"

(Gyoba, 1983; Sakurai & Gyoba, 1985). It is no accident that the same reversal also holds for the dungeon illusion (Bressan, 2006b; Bressan & Kramer, 2007), another lightness effect that, in DAT, is based on conflicting frameworks.

Can DAT explain why the black room in Gilchrist and Jacobsen's (1984) experiment did not appear white?

Gilchrist and Jacobsen (1984) created two small rooms, furnished with identical arrangements of objects, in which every surface was entirely painted black (first room) or white (second room). The black room was illuminated laterally by a strong lamp; the white room was illuminated either by a strong lamp or by a much dimmer one, such that every point had a lower luminance than the corresponding point in the black room. Observers viewed these rooms by looking through an aperture, and indicated the lightness of different test spots by selecting a matching chip from a Munsell chart, housed in a separate box and shown under a pre-set medium level of illumination (Gilchrist, 2006). Next, they were asked to adjust the illumination level on that Munsell chip to match the perceived illumination on the test spot, by looking back and forth between the room and the matching chamber.

Howe et al. take issue with the fact that the brightest spot in the black room was matched to middle gray, rather than to white as anchoring theories should, they think, predict. Following the explanation I offered in my review of a previous version of their critique, Howe et al. acknowledge that the miniature room could have been peripherally anchored to the white of the Munsell chart, therefore appearing gray rather than white. However, they choose to reject this account with the argument that the luminance of a scene viewed earlier can affect the lightness of a scene viewed later "only when there are regions common to both scenes (Annan & Gilchrist, 2004)".

This is simply not the case. What Annan and Gilchrist (2004) found is that, the larger the number of patches that remain constant in luminance and are continuously visible, the slower the new anchor is applied. But temporal anchoring does not require the presence of constant patches at all, as shown by Cataliotti and Bonato (2003). In these experiments (in which low- and high-level aftereffects were ruled out), a dim target disk was perceived as significantly darker when preceded by a bright, as opposed to a dim, disk; the bright disk acted as anchor for as long as 10 seconds of temporal separation. The darkening was even stronger when the bright anchor disk was subdivided into 6 patches ranging from black to white, showing that temporal frameworks work very much like spatial frameworks do (i.e., their strength is affected by their articulation: see Gilchrist et al., 1999; Bressan, 2006b).

A Munsell chart (16 patches ranging from black to white, on a large white surround) is of course an excellent anchor. Provided that the target is judged *after* the appearance of the anchor, it makes no theoretical difference whether the anchor precedes or follows the target in time (a difference further blurred by the fact that observers looked back and forth between the room and the matching chamber). As a final touch, it is interesting to remark that a black disk preceded by a black-to-white articulated anchor appears middle gray (mean reflectance = 31%; Cataliotti & Bonato, 2003); the brightest spot in the black room followed by a black-to-white articulated anchor (the Munsell chart) *also* appeared middle gray (mean reflectance = 28%; Gilchrist & Jacobsen, 1984, Table 1).

One may marvel at the possibility of these temporal frameworks, but frameworks *must* be fundamentally temporal for creatures that constantly move their eyes. In our experiments and in the real world, it would be difficult to find a target whose nonlocal frameworks are purely spatial. Whenever observers shift their gaze (or their attention) back and forth between a target and a matching chart, or an adjustable patch, a framework is established that is inexorably temporal in nature. If a part of our measuring instrument is brighter than the target, then *that* region is the peripheral highest luminance.[1]

As a backup argument, Howe et al. claim that, *if* the miniature rooms were indeed partially anchored to the Munsell chart, *then* the dimly lit, lower-luminance white room would have appeared darker—i.e., would have been matched to a darker shade of gray—than the brightly lit, higher-luminance black room. The logic is not watertight. Suppose the rooms were *not* anchored to the chart: how would Howe et al. explain why the low-luminance room appeared lighter than the high-luminance one?

The answer is in Gilchrist & Jacobsen (1984): the comparison between the black and white rooms was confounded by the fact that, due to a different amount of indirect illumination, the rooms contained extra information (conveyed by shadows, gradients, luminance variations) about their respective reflectances. If Gilchrist and Jacobsen are right, the black room would actually have appeared white in the absence of objects, edges, and

Munsell charts. Such a room would look a bit like an empty hemispherical dome, painted entirely black and illuminated uniformly. In Alan Gilchrist's laboratory, observers who place their heads inside such constructions do indeed report an off-white color that becomes brighter over time—that is, as the effects of prior anchors wear off.

Howe et al. seal their critique by stating that the dimly lit white room could not be anchored to the Munsell chart anyway because it "looked white". But this is not true. The median Munsell match for the dimly lit white room (the same across all test spots) was 7.5, corresponding to a reflectance of 50%, which is light gray. This squares finely with the notion of peripheral anchoring to the white of the Munsell chart. The white room would have looked white only if the Munsell chart had *not* been the highest luminance. Can this prediction be tested? It can: the brightest test spot in the brightly lit white room had a higher luminance than the white of the Munsell chart. This spot should have looked white, and it did: the median Munsell match for it was 9.5 (Gilchrist & Jacobsen, 1984, Table 1).

In summary, Howe et al. have failed to come up with even one clear challenge to the double-anchoring theory. Considering that the theory itself is still young and unrefined, this seems a sign that we are on a promising track.

---

[1] My original article contains an explicit reference to a similar situation: "In the experiment by Bonato and Gilchrist (1999), reflectance matches were made to a separately illuminated Munsell scale on white, where white had a reflectance of 90% and a luminance of 539 cd/m^2. Therefore, this white on the Munsell scale was the peripheral highest luminance" (Bressan, 2006b, p. 552).

References

Agostini, T., & Proffitt, D. R. (1993). Perceptual organization evokes simultaneous lightness contrast. *Perception, 22,* 204-272.

Annan, V., Jr., & Gilchrist, A. (2004). Lightness depends on immediately prior experience. *Perception & Psychophysics, 66*, 943-952.

Bressan, P. (2001). Explaining lightness illusions. *Perception, 30,* 1031-1046.

Bressan, P. (2006a). Inhomogeneous surrounds, conflicting frameworks, and the double-anchoring theory of lightness. *Psychonomic Bulletin & Review, 13,* 22-32.

Bressan, P. (2006b). The place of white in a world of grays: a double-anchoring theory of lightness perception. *Psychological Review, 113,* 526-553.

Bressan, P., & Kramer, P. (2007). *Remote effects on lightness.* Manuscript submitted for publication.

Cataliotti, J., & Bonato, F. (2003). Spatial and temporal lightness anchoring. *Visual Cognition*, *10,* 621-635.

Foley, J. M., & McCourt, M. E. (1985). Visual grating induction. *Journal of the Optical Society of America A, 2,* 1220-1230.

Gilchrist, A. (2006). *Seeing black and white.* New York: Oxford University Press.

Gilchrist, A., & Jacobsen, A. (1984). Perception of lightness and illumination in a world of one reflectance. *Perception, 13*, 5-19.

Gilchrist, A., Kossyfidis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X. et al. (1999). An anchoring theory of lightness perception. *Psychological Review, 106,* 795-834.

Gyoba, J. (1983). Stationary phantoms: A completion effect without motion and flicker. *Vision Research, 23,* 205-211.

Howe, P. D. L., Sagreiya, H., Curtis, D. L., Zheng, C., & Livingstone, M. S. (2007). The double-anchoring theory of lightness perception: A comment on Bressan (2006). *Psychological Review, 114,* 1105-1110.

Kingdom, F., & Moulden, B. (1991). White’s effect and assimilation. *Vision Research, 31,* 151-159.

Metzger, W. (2006). *Laws of Seeing* (L. Spillmann, S. Lehar, M. Stromeyer, & M. Wertheimer, Trans.). Cambridge, MA: MIT Press. (Original work published 1936)

Ripamonti, C., & Gerbino, W. (2001). Classical and inverted White’s effects. *Perception, 30,* 467-488.

Sakurai, K. & Gyoba, J. (1985). Optimal occluder luminance for seeing stationary visual phantoms. *Vision Research, 25,* 1735-1740.

Spehar, B., Gilchrist, A., & Arend, L. (1995). The critical role of luminance relations in White’s effect and grating induction. *Vision Research, 35*, 2603-2614.

Thayer, A. H. (1896). The law which underlies protective coloration. *Auk, 13,* 124-129.

Zaidi, Q. (1989). Local and distal factors in visual grating induction. *Vision Research, 29,* 691-697.

Paola Bressan, Dipartimento di Psicologia Generale, Università di Padova, Via Venezia 8, 35131 Padova, Italy.

I am grateful to Peter Kramer, Alan Gilchrist, Stefano Vezzani, and Branka Spehar for constructive discussion.

Correspondence concerning this article should be addressed to Paola Bressan, Dipartimento di Psicologia Generale, Università di Padova, Via Venezia 8, 35131 Padova, Italy. E-mail: paola.bressan@unipd.it.

Postscript: The Prejudice Against Frameworks

Paola Bressan

Università di Padova

In their postscript, Howe et al. (2007) raise six new miscellaneous points, which I will address in turn. First, contrary to Howe et al.'s claim that this is a post-hoc addition, frameworks were defined as "regions that are likely to share the same illumination" all along (e.g. Bressan, 2006b, p. 530): in fact, the whole theory is constructed around this notion (see pp. 529-531, 544-546, and 550). The gist is that regions that share some photometric and spatial constraints (the grouping principles) are likely to lie within the same shadow or splash of light; weighted averaging of the lightnesses computed inside and outside such local frameworks helps to discount illumination and recover reflectance. Second, contrary to Howe et al.'s allegations of circular reasoning, no preliminary estimation of either lightness or illumination is of course necessary to determine the frameworks: we exploit dumb grouping principles instead. Third, DAT does indeed predict that regions near the border receive more contrast than regions farther away, but contrary to Howe et al.'s contention this prediction is confirmed by the ubiquitous phenomenon of edge enhancement (e.g., Remole, 1977). Fourth, when the luminance hierarchy is overturned in grating induction, the relative weights of the two frameworks (which also depend on their size and hence on the spatial frequencies of the regions concerned) are crucial for understanding "whether the regular effect will simply diminish, disappear, or reverse" (Bressan, 2007). The work by Kingdom, McCourt, and Blakeslee (1997), rather than disconfirming the theory, supplies an example of the first case; the work by Spehar, Gilchrist, and Arend (1995) of the second; the work by Sakurai and Gyoba (1985) of the third. Fifth, the information about illumination, including gradients and shadows, far from being "an entirely new concept" is amply discussed in the original theory, and lies at the heart of the idea of *overlay frameworks* (pp. 530 and 544-545).

Sixth and last, in view of the four appendices and eight graphs showing predicted data points against observed ones in Bressan (2006), and considering that such predictions stem from a plain equation that everyone can try for themselves, it is peculiar that my theory could be described as “less precise” than the original anchoring theory (which did not put forth quantitative predictions). This criticism echoes a misrepresentation of frameworks as post-hoc exercises of data fitting. Yet, the model uses only two relative weights, both theoretically constrained. Together, they reflect the importance of the local surround (weighing, respectively, its information content and its information reliability) relative to the rest of the scene. These weights are chosen with the restriction that each of them must be the same for *all* data points (see the Appendices in Bressan, 2006), a massive constraint. The double-anchoring theory of lightness is fully falsifiable and in all likelihood false; but still closer to the truth than its predecessors.

References

Bressan, P. (2006). The place of white in a world of grays: a double-anchoring theory of lightness perception. *Psychological Review, 113,* 526-553.

Bressan, P. (2007). Dungeons, gratings, and black rooms: a defense of double-anchoring theory and a reply to Howe et al. *Psychological Review, 114*, 1111-1115.

Howe, P. D. L., Sagreiya, H., Curtis, D. L., Zheng, C., & Livingstone, M. S. (2007). The double-anchoring theory of lightness perception: a comment on Bressan (2006). *Psychological Review, 114,* 1105-1110.

Kingdom, F. A., McCourt, M. E., & Blakeslee, B. (1997). In defence of “lateral inhibition” as the underlying cause of induced brightness phenomena: a reply to Spehar, Gilchrist and Arend. *Vision Research, 37*, 1039-1044.

Remole, A. (1977). Brightness enhancement versus darkness enhancement at a border. *Vision Research, 17,* 1095-1100.

Sakurai, K. & Gyoba, J. (1985). Optimal occluder luminance for seeing stationary visual phantoms. *Vision Research, 25,* 1735-1740.

Spehar, B., Gilchrist, A., & Arend, L. (1995). The critical role of luminance relations in White’s effect and grating induction. *Vision Research, 35*, 2603-2614.