Bressan, P. (2006). The place of white in a world of grays: a double-anchoring theory of lightness perception. *Psychological Review, 113,* 526-553. doi: 10.1037/0033-295X.113.3.526

The Place of White in a World of Greys: A Double-Anchoring Theory of Lightness Perception

Paola Bressan
Università di Padova[1]

The specific grey shades in a visual scene can be derived from relative luminance values only when an anchoring rule is followed. The double-anchoring theory I propose in this paper, as a development of the anchoring theory of Gilchrist et al. (1999), assumes that any given region (a) belongs to one or more frameworks, created by Gestalt grouping principles, and (b) is independently anchored, within each framework, to both the highest luminance and the surround luminance. The region's final lightness is a weighted average of the values computed, relative to both anchors, in all frameworks. The new model not only accounts for all lightness illusions that are qualitatively explained by the anchoring theory but also for a number of additional effects, and it does so quantitatively, with the support of mathematical simulations.

Keywords: lightness illusions, simultaneous lightness contrast, highest luminance, anchoring theories, illumination discounting

> All at once his white shirt blazed out, and I came out after him from shadow into full sunlight...
>
> —Ursula LeGuin, *The Left Hand of Darkness*

Our visual world can be treated as a collection of groups of surfaces that belong together, either by design, like the black and white stripes on the zebra's back or the four regions in Figure 1a, or by accident, like the combination of sky fragments and tree branches through my window every morning at eight o'clock. Groups such as these have been called frameworks (Gilchrist et al., 1999). Depending on the structural complexity of its context, a region can simultaneously belong to two or more nested frameworks, as is indeed the case with most objects in our visual experience. In the example of Figure 1a, what we may call the global

---

[1] This article is a result of the cognitive dissonance stemming from my enchantment with the anchoring theory and my subsequent, and unplanned, experience of double-increment illusions. I am indebted to Peter Kramer, Dejan Todorović, Dave Rose, Ennio Mingolla, Luigi Burigana, Joe Cataliotti, Fred Bonato, Nicola Bruno, Tom Troscianko, Elias Economou, Suncica Zdravkovic, Piers Howe, Daniele Zavagno, and Fred Kingdom for their critical reading of one or the other of the several drafts of this manuscript. Special thanks to Alan Gilchrist for innumerable exchanges—an honor and a pleasure.

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

framework (by temporarily ignoring the page and the visible rest of the room) is composed of two side-by-side local frameworks, each containing a grey patch on a uniform background. The interesting aspect of this familiar display is that the two patches are the same physical shade of grey, but the one in the black field appears lighter than the other. In spite of its apparent irrelevance to the issue of lightness, the multiple-framework idea, developed by Gilchrist and his collaborators, explains precisely why this should be so.

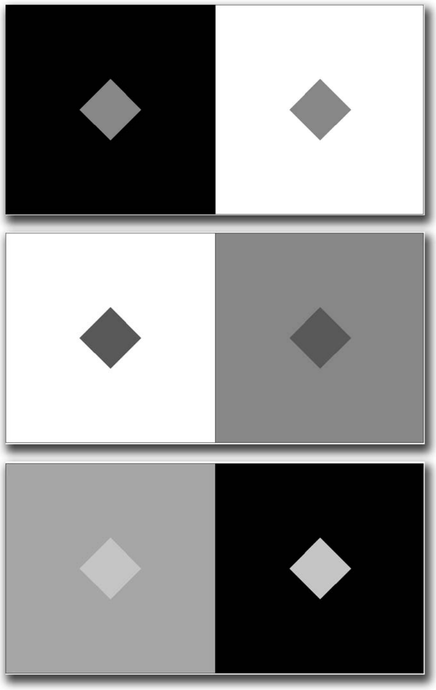

Figure 1. Simultaneous lightness contrast with one increment and one decrement (A, top panel), with two decrements (B, middle panel), and with two increments (C, bottom panel).

The model says that the lightness of any given surface is a weighted average of the lightness of the surface when anchored to its local framework and the lightness of the surface when anchored to the global framework. These frameworks are nested, that is the local framework is included in the global. Within each framework, the role of anchor is assigned to the highest luminance, which is locally given a value of white. A corollary rule applies to the case in which the darker region takes up more than half of the visual field: in this case, such a region tends to lighten, and the larger it becomes the lighter it appears.

Why do we need to assume that the apparently effortless judgment of the lightness of a grey patch must be the result of such complex computations? We do because a region's lightness (that is, its perceived achromatic color) does not correspond in any direct way to its luminance (the amount of light it emits, or reflects). For that matter, its luminance does not correspond in any direct way to its physical grey shade either. Any given luminance value conflates the contributions not only of the actual grey shade of a surface (which can vary by a factor of thirty to one), but also, and especially, of the intensity of the light that shines on it (which can vary by a factor of a *billion* to one). Because these two allotments can never be torn apart, the information carried by luminance is hopelessly uncertain. Such local uncertainty cannot be reduced other than by taking in a wider portion of the visual scene (and with it, more information about the prevailing illumination conditions) when assessing lightness. This explains why any given luminance value can be perceived as virtually any shade of grey—all the way from black to white—depending on its context (Gelb, 1929).

Taking the context into account can become computationally simple if the starting point is the estimate not of luminance *values*, but of luminance *ratios* between parts of the scene (Wallach, 1948). Yet luminance ratios are as ambiguous as everything else, because they can only produce lightness ratios (of the bizarre form, this figure is three times as light as its background) rather than the familiar absolute lightness values (of the form, this figure is dark grey and its background is black). This is where an anchoring rule must be called into play. An anchoring rule establishes a point of contact between relative luminance values and the black-to-white scale of lightnesses. In Gilchrist et al.'s (1999) model, this point of contact joins the *highest luminance* in the scene and the value of *white* on the lightness scale.

In an achromatic world, then, the region with the highest luminance will be unambiguously seen as white, and all other regions will be perceived as shades of grey, depending on their luminance ratio to such white. In the global framework of Figure 1a, the big light area on the right works as an anchor, and the two small patches are assigned identical grey values relative to it. In the local frameworks, however, the lightness assignments are different for the two targets: grey for the one on the right (which gets compared to the same anchor), white for the one on the left, because, being the highest luminance of its local framework, it works as the local anchor. Thus the target on the light background is globally grey and locally grey, whereas the target on the black background is globally grey and locally white. When the local and global values are combined in a weighted average, the latter target will yield a perceptually lighter grey.

Suppose now that both targets are luminance increments relative to their backgrounds, as in Figure 1c. The anchoring model requires that this display show no contrast at all (Gilchrist et al., 1999). The two incremental targets have identical lightness assignments not only globally but also locally, because each is the highest luminance (and thus, white) within its local framework.[2]

This prediction is wrong, though, as shown by Figure 2. The four horizontal rows of diamonds are physically the same white as the page. Being the highest luminance in the whole image, they necessarily represent the highest luminance in their local frameworks, no matter how such frameworks are defined. The diamonds should thus have identical lightness assignments both globally and locally; yet, those that represent a larger increment relative to the luminance that immediately surrounds them look whiter than those that represent a smaller increment.

---

[2] Consistent with such expectation, a few studies (e.g., Arend & Spehar, 1993; Gilchrist, 1988; Heinemann, 1955) found no contrast effect with double increments. On the basis of their own data, Bressan and Actis- Grosso (2001) have argued that such a failure was due to the choice of luminances and, more specifically, to a surround luminance that was too high (as in Gilchrist, 1988) or to a target luminance that was too low (as in Arend & Spehar, 1993, and Heinemann, 1955). Indeed, Bressan and Actis-Grosso (2001) found no significant double-increment effects for the luminances used in previous works.

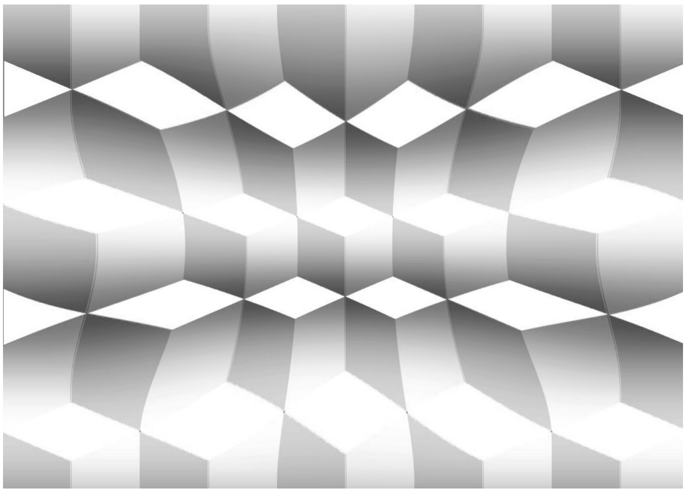

Figure 2. The Christmas wall-of-blocks, that is, a wall-of-blocks as after a snowfall. The tops of all blocks (four horizontal rows of diamonds) are equally white. This is my version of Logvinenko's (1999) figure, which is, in turn, a modified version of Adelson's (1993) "wall-of-blocks" where the sides of the blocks are not uniform but filled with a luminance gradient. In the version I introduce here, the tops of all blocks represent increments. The illusion is all the more impressive in that the actual luminance differences between the ends of each gradient can be fairly subtle.

The Christmas wall-of-blocks points to a problem with the anchoring theory's basic tenet that the lightest region can have no lightness value other than white. The theory acknowledges that the lightest region can take on additional qualities, such as luminosity, but only as a consequence of an increase in the area of the dark surround. This idea is expressed more formally with the so-called "area rule": the darker area in a simple framework, if it is also the larger, moves towards white, pushing the lighter region toward luminosity (Gilchrist et al., 1999; Li & Gilchrist, 1999). Yet, highest-luminance regions can look off-white, white, fluorescent, or even luminous *regardless* of the areas of their surrounds (Bressan & Actis-Grosso, 2001), as shown by the mutable appearance of the moon in the evening sky (Bressan, 2005).

The variant of the anchoring theory I describe in this paper originated as an attempt to reconcile the principle of anchoring to the highest luminance with the existence of double-increment illusions. The new model includes an additional anchor, but is thereby able, first, to explain *all* the data explained by the original model with fewer ad hoc assumptions (e.g., dispensing with concepts such as scale normalization or insulation); second, to account for a number of empirical findings that the original model left unexplained. Examples of the latter are simultaneous contrast with incremental targets; several data on remote and reverse contrast, bipartite domes, equivalent backgrounds, and the Staircase Gelb effect; hypercompression; insulation; lightness enhancement in subjective figures; lightness of objects under two separate illumination levels; the influence of depicted illumination on lightness; modulation of lightness by attention; and individual differences in lightness assessment.

In the double-anchoring model, the first anchor is the highest luminance, the second is the *surround* luminance. I shall argue that anchoring to the highest luminance evolved for interpreting luminance changes due to surface colors, and anchoring to the surround luminance evolved for interpreting luminance changes due to light sources.

Essence of Double Anchoring

The new model proposes that, within each framework, the lightness of the target region is determined not only by its luminance ratio to the highest luminance (HL step), but also by its luminance ratio to the surround luminance (surround step). Because they are anchors, highest luminance and surround luminance are defined as white. The weighted average of these two luminance ratios is the "territorial lightness" of that region in that framework (of course the term "lightness" is used improperly here, since it does not coincide with a percept). The final lightness value of the target region is the weighted average of all its territorial lightnesses. A surround-as-white rule was advocated by Gilchrist and Bonato in 1995, and eventually repudiated on account of its alleged failure to capture several important data (Bonato & Gilchrist, 1999; Li & Gilchrist, 1999); as we shall see, the reinterpretation of this rule in the double-anchoring model is perfectly consistent with those data.

In the double-anchoring model, the concept of "surround" departs a little from its intuitive meaning. The surround is defined as any region that groups with the figure. For this reason, a surround need not be retinally adjacent to the figure. For example, each window on the front of a building belongs to two separate frameworks. One is founded on the Gestalt principle of proximity, in its strong form of adjacency, and consists of the window plus the face of the building. In this framework the face of the building (the wall) serves as surround. The other is sustained by the Gestalt principle of shape similarity, and consists of the window plus all the other windows: in this framework, it is these windows that serve as surround. Whether the color of these other windows visibly affects the color of the target window, will depend on the balance between the grouping forces at play in the two co-existing frameworks (see Figure 9). This example makes it clear that the condition of being a figure is a contingent, and not an inherent, property of a region. The window that plays figure at this instant in time will serve as surround when I move my gaze towards a different window, an instant from now.

In the natural world, any given framework may be thought of as composed of figure and surround. But in special (impoverished) conditions, the visual scene may comprise one or more surrounds only. This occurs when we look at the empty sky, for example: a single framework consisting of just one surround and no figure. In Gilchrist's laboratory, the same happens when subjects place their heads inside a large illuminated hemispherical dome uniformly painted black, or grey. When a sector of the dome is painted a different shade of grey, the dome turns into a single framework composed of two surrounds, as we shall see later. Any region that has no fully delimited boundaries (as an empty sky), is always, in the language of the model, a surround.

Frameworks: What They Are and What They Mean

Which are the relevant frameworks in a scene? Usually we need to consider only two: local and peripheral. The local framework of a figure consists of the figure and its immediate surround. The peripheral framework links the figure to the rest of the visual field. In the double-anchoring model these two frameworks are not nested: hence, the peripheral framework does not contain the luminance of the local surround. This distinction between local framework and *peripheral* framework (rather than *global*, as in the original anchoring model) is similar to the one made by Kardos (1934) between "relevant" and "foreign" fields. Kardos was the first to suggest that the lightness of a target is not only determined by the target's own framework, but also by an external framework.

In case of complex displays, we occasionally need to take into account additional, intermediate frameworks between the local and peripheral ones (called *super-local*). In general, we can ignore assignments for all frameworks higher than local, on the grounds that they will simply dilute, but never reverse, the effects generated in the local frameworks. But in some displays, targets can receive identical values in their local frameworks and different values in super-local frameworks, and in this case the latter assignments need to be taken into account. An especially clear example is remote contrast, that will be discussed later.

The frameworks to which a target belongs are determined by the spatial and photometric grouping factors that link the target to the other regions in the scene. Examples of spatial factors are *proximity* (e.g. Ben-Av & Sagi, 1995), whose strong form is adjacency; *common region* (Palmer, 1992), that is the tendency of elements located within the same closed region of space to group together; *alignment* or "good continuation", whose strong form is the T-junction (e.g. Todorovic´, 1997); *depth similarity* (e.g. Gogel & Mershon, 1969), whose strong form is coplanarity; and *shape similarity* (e.g. Ben-Av & Sagi, 1995).

Examples of photometric factors are luminance polarity and similarity (e.g. Masin, 2003a). *Luminance polarity* means that, other grouping forces being equal, grouping will tend to occur preferentially between regions with the same contrast sign. *Luminance similarity* means that, other grouping forces being equal, grouping will tend to occur preferentially with the region, or regions, whose luminance is closer to the target's. When luminance polarity and luminance similarity are pitted against each other, grouping is predominantly affected by either the former or the latter for some observers, but equally affected by both for other ones, as shown very clearly by Masin (2003a). When luminance polarity and luminance similarity concur, the resulting groups can be remarkably strong. Different regions sharing a common illumination, either in light or in shade, are the most notable example of such groups.

Frameworks created by "hard" grouping principles such as those listed above tend to behave as stable entities, and are little or not at all affected by voluntary control. But frameworks can also be sustained by "soft" grouping principles, such as attention or past experience; or can be based on a compelling grouping force, say adjacency, but weakened by another grouping force pushing in the opposite direction, say binocular disparity values indicating separation in depth. In the context of lightness experiments, unstable or conflicting frameworks are likely to reveal inter-individual differences (due for example to variations in attention, fixation patterns, experience, or interpretation of the task demands), and potentially also intra-individual differences, under the form of variability across stimulus repetitions or sessions (Bressan, in press).

It should be clear from the above that, in the model, frameworks are only a convenient way of expressing the idea that each point in a scene is influenced by every other point, and that how strong this influence is hinges on some affinity measure between the points—that is, on how strongly they "group" with each other. Other things being equal, for example, nearby points matter more than distant points (proximity), and same-disparity points matter more than different-disparity points (coplanarity). What we do here is aggregate all identical points into a "region" and treat that as a unity. The concept of frameworks is thereby not really necessary, but it makes life simpler.

Overlay Frameworks

It has been repeatedly argued (e.g. Logvinenko & Ross, 2005) that some illusions must be due to a misjudgement of apparent illumination, and cannot possibly be handled by anchoring theories. Such positions are based on the erroneous notion that frameworks function solely as adjoining pieces of a flat mosaic, a misunderstanding that stems from the assumption that grouping processes operate only at the level of the retinal image. But all the available evidence (see Palmer, Brooks & Nelson, 2003, for a thorough review) shows that grouping occurs at least twice, once before and once after depth and constancy processing have been completed. "Late" or top-down grouping alters the content of the visual representation according to attributes typical of 3-D objects and layouts, rather than of 2-D patches.

An especially instructive case of late grouping concerns what I will call *overlay frameworks*. Two adjacent frameworks give rise to an overlay framework when some grouping factor (usually photometric, such as luminance polarity or similarity) joins regions *within* them, while some other grouping factor (typically, good continuation) joins regions *across* them. Overlay frameworks are thus the superordinate level of a nested hierarchy. They give rise to the visual impression of overlapping layers, of which the closer appears as transparent, and the farther appears as seen through it. I call this grouping "late" in the sense that it builds on the "early" grouping of patches into standard frameworks, and that it goes back to update the lightness assessment of each patch by splitting it into two separate values. The first is the lightness of the patch at the farther depth plane (the surface seen through the overlay), the second is the lightness of the patch at the closer depth plane (the transparent overlay).

In its strong form this type of grouping is embodied by X-junctions (see Adelson, 2000), but junctions are indispensable only to the extent that they are needed for grouping purposes. If regions lying in two adjacent standard frameworks group together anyway (for example, by alignment or shape similarity), an overlay framework can be generated in the absence of junctions. In the natural world, transparent or translucent media (such as unclean waters, haze, or smoke) and sharp or gradual variations in illumination (such as spotlights, shadows, or lighting ramps) give rise to overlay frameworks.

Overlay frameworks have two important effects on their component frameworks: they make them especially stable (elements are said to form stable groups when they are difficult or impossible to ungroup, or unavailable for alternative groupings), and they make them especially strong (weightier in the final average). The latter property stems from the former, and reflects the general principle that regions grouping with the target in a stable, dependable manner (say, by common illumination) influence it more than regions grouping with it ephemerally and unreliably (say, by contiguous eye fixations).

Overlay frameworks normally arise from adjoining frameworks, but can in their limiting case be founded on a single simple framework. A familiar example is a cast shadow on an otherwise homogeneous area, say a snow field. The role of top-down grouping in creating overlay frameworks where grouping forces other than proximity are missing is obvious. At the retinal level there is, of course, no information as to whether the less luminant region is part of an overlay framework (e.g., it is a shadow) or of a standard framework (e.g., it is a patch of dirty snow). Upon creation of a standard framework, such a region would be assessed relative to the white snow (highest luminance and surround) and matched to dark grey. Upon creation of an overlay framework, however, the same region receives two separate lightness assessments, one at each depth plane, returning a sense of white snow under a shadow.

Rationale of Frameworks

The fundamental assumption of an anchoring model is that recovering the color of objects does not require complex inferences about the illuminant. We only need to assess the target luminance relative to some reference luminance of "known" lightness: say that we choose as such anchor the highest luminance, and call it "white". In a flat, uniformly lit world this would be enough; but because the amount of light that hits them varies, regions must be anchored within their local illumination level also (if they were not, our estimates would be grossly mistaken; for example, objects in the shadow may appear black). I suggest that such local gauging is accomplished by grouping the target with regions that are likely to share the same illumination, and anchoring it to the luminance that, *within* the group, is labeled as "white". To parse the scene into groups, we do not need to know about the illuminant, but we use a set of crude rules: proximity, depth similarity, alignment, luminance polarity and similarity, and so on. In nature, regions that share these properties are likely to be equally illuminated.

Yet the final estimate needs to be global, not local; if not, the lightness of objects would be inescapably linked to their groups' local "whites"—with the result that, for example, an object that moves onto a background darker than itself would actually turn white. What we do, then, is assess the target relative to both the *local* "white" (which yields an estimate of the target's lightness under the local illumination) and the *peripheral* "white" (which yields an estimate of the target's lightness under the non-local illumination). We reach a global estimate of the "true" color of the target simply by averaging these two territorial estimates, weighted by their reliability. In the language of the model, these two groups are the local framework and the peripheral framework.

Formal Rules of Double Anchoring

Double Anchoring and Weighting Within Frameworks

In single-framework images, the lightness of a target (defined as its matching luminance on a white surround) is the result of the independent application of the surround and highest-luminance rules. Surround and highest luminance are regarded as white, and the appearance of the target depends on its relationship to both regions, according to the formula

$$L_M = (L_t/L_s + L_t/L_h) \times L_w, \qquad (1)$$

where L_M is the predicted matching luminance on white, L_t is the luminance of the target, L_s is the luminance of the surround, L_h is the highest luminance in the framework, and L_w is the luminance of white. The luminance of (the region that is treated as) white is not fixed, of course, but depends on the contextual luminances. Usually it is the highest luminance in the visual field.

The resulting lightness L_M, or any of the intermediate "lightnesses", can be larger than the lightness of white. We will call this *superwhite*. Superwhite values should always be read as if placed on an ordinal, rather than an interval, scale.

The two ratios of Equation 1, of course, need to be appropriately weighted, in order to express the relative importance of the surround and HL steps. The final lightness value of a target in a simple image is therefore the weighted arithmetic mean of the values computed at the two steps, that is

$$L_M = [(L_t / L_s \text{ x } W_s + L_t / L_h \text{ x } W_h)/(W_s + W_h)] \text{ x } L_w, \qquad (2)$$

where W_s is the weight of the surround step and W_h is the weight of the HL step. A number of factors can affect this balance. The relative weight of the surround step is, most notably, a direct function of (a) the size of the surround relative to the target (*weight/area* rule); (b) the articulation of the surround, that is the number of different regions it contains (*weight/articulation* rule); and (c) the absolute luminance of the surround (*weight/luminance* rule*)*. In simpler words, a large, articulated, and well-lighted surround makes for a heavier anchor—reflecting the fact that, in nature, it makes for a context that is richer in information, and thus more reliable.

Weighting Between Frameworks

The final lightness of a target in a scene containing two frameworks, local and peripheral, can be expressed as:

$$L_M = [(T_l \text{ x } W_l + T_p \text{ x } W_p)/(W_l + W_p)] \text{ x } L_w, \qquad (3)$$

where W_l and W_p are the weights given to the local and peripheral frameworks, and T_l and T_p are the territorial lightnesses in the local and peripheral frameworks. For simplicity (given this is mathematically equivalent), each territorial lightness is determined here by applying Equation 2 without the multiplication by L_w, and this multiplication is done at the final computation stage instead.

The strength of a framework (operationally, the weight assigned to it) is 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.

Double anchoring within each framework is a necessity only in principle. For the practical purpose of data modeling, it is indispensable in local frameworks, but not in the peripheral framework, where the contribution of surround anchoring is typically so small as to be negligible. Hence, in all calculations presented in this paper, double anchoring will be applied in local and super-local frameworks, and simple HL anchoring in peripheral frameworks.

Testing the Model: Simultaneous Lightness Contrast

The double-anchoring model will now be tested by pitting its predictions against those of the original anchoring model. This will be done not only for the anchoring model's predictive failures, such as simultaneous contrast with double increments, but also and especially for its predictive successes. Whenever quantitative data are available, I shall rely on the luminance or reflectance values used in the experiments and give precise predictions (full calculations are provided in the Appendixes), which can be compared with the actual results. If the data to be modeled are expressed as reflectances, rather than as luminances, I will simply replace luminance values with reflectance values, and calculate R_M (matching reflectance on white) rather than L_M. There was no quantitative modeling in Gilchrist et al. (1999); this is the first attempt to do so within an anchoring theory.

In all the simulations presented in this paper, the predicted values are obtained by choosing the optimal weight for the local framework relative to the peripheral framework, and the optimal weight of local surround anchoring relative to local HL anchoring, to allow a fit as close as possible to the observed data. Although the fine tuning is done by adjustment, the choice of weights is not arbitrary, but always consistent with the weighting rules illustrated earlier, as we shall see. Of course the *same* weights are used for all points, which produces one single slope (and not just *any* slope).[3]

The simplest lightness illusion is called simultaneous lightness contrast, and comes in the three variants of Figure 1. Relative to the luminance of their backgrounds, the two physically identical target patches can be one increment and one decrement as in Figure 1a, two decrements as in Figure 1b, or two increments as in Figure 1c. From a double-anchoring stance, the effects observed in the three types of display are of a different nature and strength. Contrast effects are predicted to be largest with one increment and one decrement. The decrement is anchored to its surround, which is also the highest luminance (see Figure 3, top row), whereas the increment is white at the HL step and superwhite at the surround step (Figure 3, bottom row). If the decrement sits on white (making for the largest possible contrast effect), then locally we have a grey target vs a superwhite target.

We expect a much smaller effect with two decrements: they are both anchored to their surrounds, which also have the highest luminance. So, the two targets do receive different values (at both steps), but neither target is white or superwhite. If one of the two decrements sits on white (making for the largest possible contrast effect), locally we have a grey target vs a lighter-grey target. We predict an even smaller effect with two increments: they are both white at the HL step, and stand apart at the surround step only, where they receive different values of superwhite. If one of the two increments sits on black (making for the largest possible contrast effect), then, locally we have a superwhite target vs a less-superwhite target.

[3] All simulations are available on request.

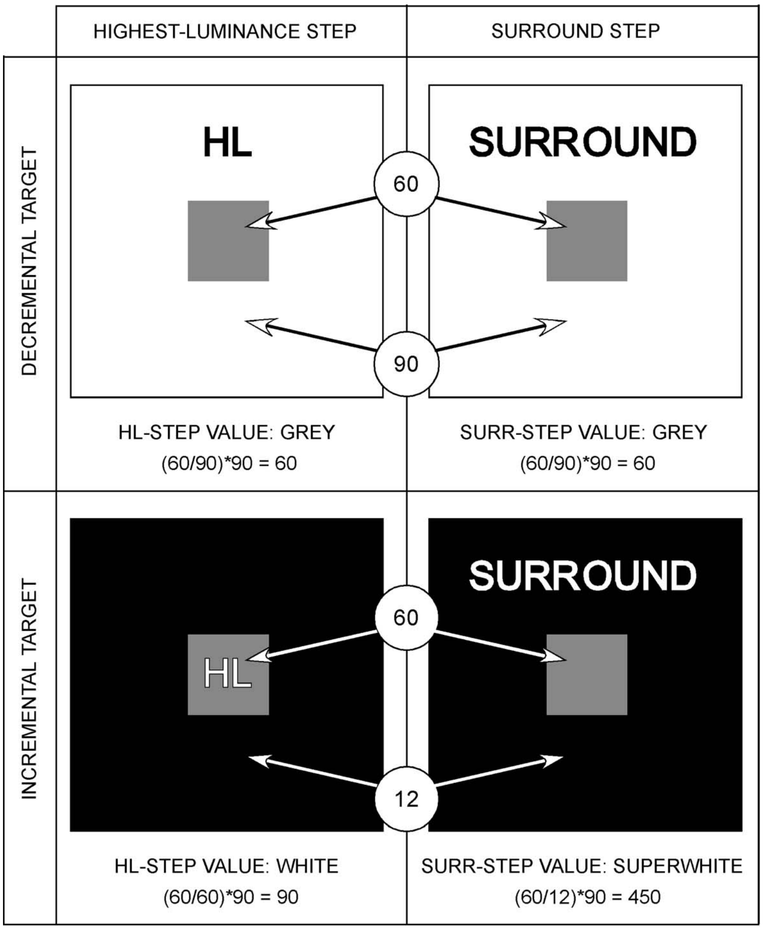

Figure 3. Double-anchoring explanation of simultaneous lightness contrast. The diagram indicates which regions serve as highest luminances (left column) and surrounds (right column) for the decremental (top row) and incremental (bottom row) targets, in the respective local frameworks. For illustrative purposes, we set the reflectances of white, gray, and black at 90%, 60%, and 12%, respectively; local values are computed for each target at each step by applying Equation 1. The value of 450% obtained for the incremental target at the surround step means that, in the absence of further anchoring, such target would be seen as luminous. The final local value for each target is a weighted average of its two values (where the HL step weighs much more than the surround step). In the peripheral framework, where only highest-luminance anchoring is expected to have a significant effect, the decremental (locally gray) target groups weakly with the black background and receives a value of white, and the incremental (locally superwhite) target groups weakly with the white background and receives a value of gray. These assignments will dilute the illusion even if we ignore the rest of the scene; for a striking illusion, each half display should be presented separately and be the only illuminated object in the visual field.

The traditional simultaneous lightness contrast display of Figure 1a is simply a combination of the strongest case for decremental targets (white surround) with the strongest case for incremental targets (black surround). We will now see in detail how the double-anchoring model predicts the effect of surround luminance on decremental and incremental targets respectively, by fitting Equation 3 to actual experimental data. The model suggests that lightness illusions occurring in local frameworks, and not helped by peripheral anchoring, can be strengthened by eliminating the peripheral frameworks. This condition can be approximated by presenting a luminous or especially-illuminated display in the dark, which does not completely eliminate the peripheral framework, but reduces its relative weight by drastically decreasing its articulation and absolute luminance. For this reason, the data modeled here were collected by presenting, on a computer screen, displays that took up most of the observer's visual field, in a darkened room entirely painted flat black.

Simultaneous Contrast with Decremental Targets

In this experiment (Guadagnucci, 2002) 50 observers used the method of adjustment to vary the luminance of a decremental test square, centered on a white background, to match the achromatic color of a decremental comparison square, centered on a variable background. (See the caption of Figure 4 for a concise description of the method.)

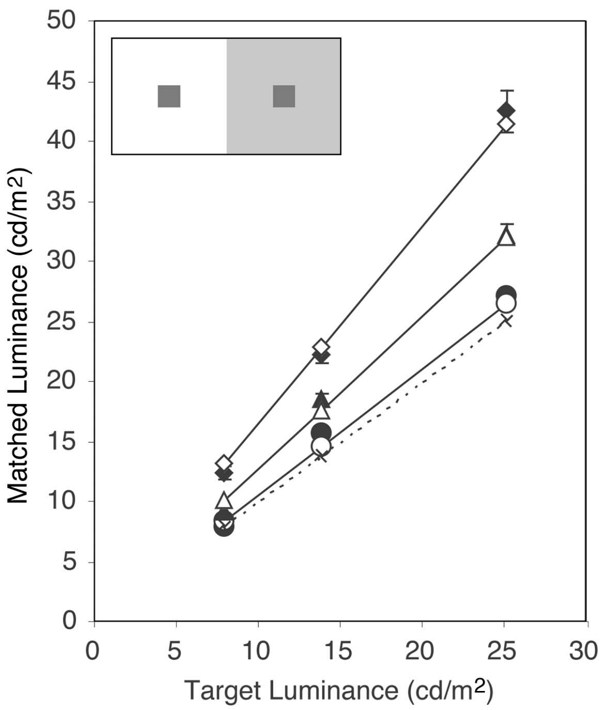

Figure 4. Matched luminance of a decremental test square on a 111.9 cd/m2 background as a function of the luminance of a decremental comparison square on a variable background, for three luminances of the variable background (top to bottom, 35.2, 57.9, and 93.2 cd/m2, respectively; Guadagnucci, 2002). The dashed line shows objective luminance matches. Solid symbols indicate the observed values, and open symbols indicate the values predicted by the double-anchoring model for a local: peripheral weighting of 3:7 (see Appendix A). Stimuli were generated by a personal computer on a 17-in. Philips Brilliance 105 screen. Each stimulus consisted of two 0.7° x 0.7° patches centered on two 16° x 22° adjacent surrounds and was shown twice, with the adjustable patch once on the right and once on the left. Fifty naive participants viewed the monitor from a distance of 60 cm in an otherwise dark room, entirely painted flat black, after adaptation to ambient light. They were asked to adjust the luminance of the patch on white to match the achromatic color of the patch on the variable background. The icon in the upper left corner represents the stimulus display.

Here, as in Figure 1b, the local framework of each square is the background on which it sits (grouping by adjacency), and the peripheral framework is the background nearby (grouping by proximity). Figure 4 shows mean test luminances as a function of comparison luminances, for three different luminances of the variable background, together with the values predicted by the double-anchoring model (Equation 3) under a local: peripheral weighting of 3:7. Note that, in the case examined here, the contrast effect is fully brought about by either HL or surround anchoring and does not need double anchoring, because a decremental target receives identical values at both steps (hence their relative weighting in this case is immaterial).

Simultaneous Contrast with Incremental Targets

Considering that incremental targets only differ at the surround step, it comes as no surprise that double-increment illusions (Figure 1c) have rarely found a place in textbooks. But the double-anchoring model suggests that we can make them obvious to the eye, or even striking, if we do one of two things. The first is increasing the weight of the local (relative to the peripheral) assignments in the final average. This can be done, for example, by presenting our displays in the dark under special illumination, or by strengthening standard frameworks via superordinate overlay frameworks, as in the Christmas wall-of-blocks of Figure 2 (where gradients signal lighting ramps). The second is increasing the relative weight of the surround step: by articulating the surround itself via fragmentation, as in Bressan and Actis-Grosso (2006), or simply by enlarging it considerably relative to the target patches, as in Bressan and Actis-Grosso (2001).

In the latter work, 20 observers used the method of adjustment to vary the luminance of an incremental test square, set on a black background, to match the achromatic color of an incremental comparison square, set on a variable background. The ratio between surround and target areas was about 700:1. The square always looked lighter on black than on any other surround. Illusion magnitude was an increasing function of surround luminance and target luminance. Both facts make sense in the model. First, for any given target, increasing surround luminance means decreasing the ratio of the target to the surround luminance, thereby making it more and more different from the ratio of the target to the black surround luminance. This strengthens the simultaneous contrast illusion. Second, for any given surround, increasing target luminance corresponds to increasing the ratio of the target to the surround without changing the ratio of the target to the highest luminance (the target itself). This results in a relatively higher value for the same target on black, and, again, in a stronger simultaneous contrast illusion.

Both points can be appreciated quantitatively by inspecting Figure 5, where some of the data reported by Bressan and Actis-Grosso (2001) are shown as percent deviations from the objective matching luminance, together with the predictions of the double-anchoring model.[4] The original anchoring model predicts no errors.

---

[4] These data are expressed as percent deviations from the objective matching luminance rather than as matched luminances, because the model's equations generate the match that would be selected for the adjustable square on a *white* background. In this experiment, the adjustable square was presented on a black background, not on a white one. Hence, although it was used as a measuring instrument, it was itself lightened by its own background. The problem is circumvented by plotting the ratio between the matched (predicted) value of the standard on the variable

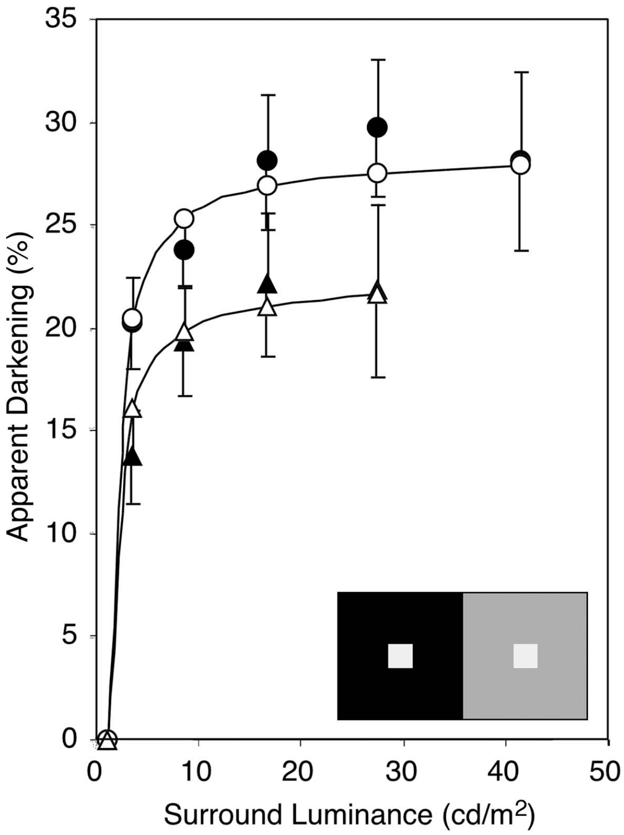

Figure 5. Apparent darkening of an incremental comparison square (of either 66.05 cd/m^2 [circles] or 47.82 cd/m^2 [triangles]) on a variable background, as measured by adjusting an incremental test square on a 0.99 cd/m^2 background. Darkening is expressed as percent deviation from the objective matching luminance; bars depict standard errors (data based on Bressan & Actis-Grosso, 2001, Experiment 1). Apparatus, stimulus sizes, and procedure were as reported in the caption of Figure 4. Participants were asked to adjust the luminance of the patch on black to match the achromatic color of the patch on the variable background; the icon in the bottom right corner depicts the stimulus display. The open symbols represent the values predicted by the double-anchoring model for a peripheral weight of zero and a surround-step:HL-step weighting of 0.006:1 (perfect fit for the last point). Peripheral anchoring can be disregarded, because the local and peripheral highest luminances are the same. The original anchoring model predicts no errors (i.e., two overlapping straight lines which in turn overlap the *x*-axis).

Although simultaneous contrast is much weaker with two increments than with one increment and one decrement, we are not talking about a small effect here. In the typical simultaneous lightness contrast display presented on paper (see Figure 1a), the mean logarithmic difference between the two targets is around 0.10 (Gilchrist, 1988). In the double-increment effect reported by Bressan & Actis-Grosso (2001) and modeled here, this difference can be as large as 0.16—that is, 50% stronger than the average textbook contrast illusion.

---

background and the matched (predicted) value of the standard on the black background, which is equal to its objective matching luminance.

Two Loci of Error: Darkening of Targets on White Backgrounds

"Locus-of-error" experiments show that, in simultaneous contrast displays such as that of Figure 1a, the target on black lightens considerably, and the target on white also darkens a little (Annan, Economou & Gilchrist, 1998; see also Gilchrist et al., 1999). The first result is easily predicted by the original anchoring model. The second has been accounted for by a "scale normalization" principle, consisting in an expansion of the range of lightnesses relative to the range of luminances. In simpler words, even when luminances only specify a bunch of greys, the perceived range of shades tends toward that between black and white. The perceived difference between the white background and the grey target increases; but since the white background, being the highest luminance, fails to become substantially whiter, the expansion is mostly expressed as a darkening of the target.

The grey target on white darkens, but relative to what? In these experiments observers matched the target square to one chip of a scale of grey chips sitting on a black-and-white checkerboard, such that each chip bordered black and white equally. Now, the fact that a target on white looks darker than a physically equal target on a checkerboard comes as a natural consequence of double anchoring, as shown in Figure 6.

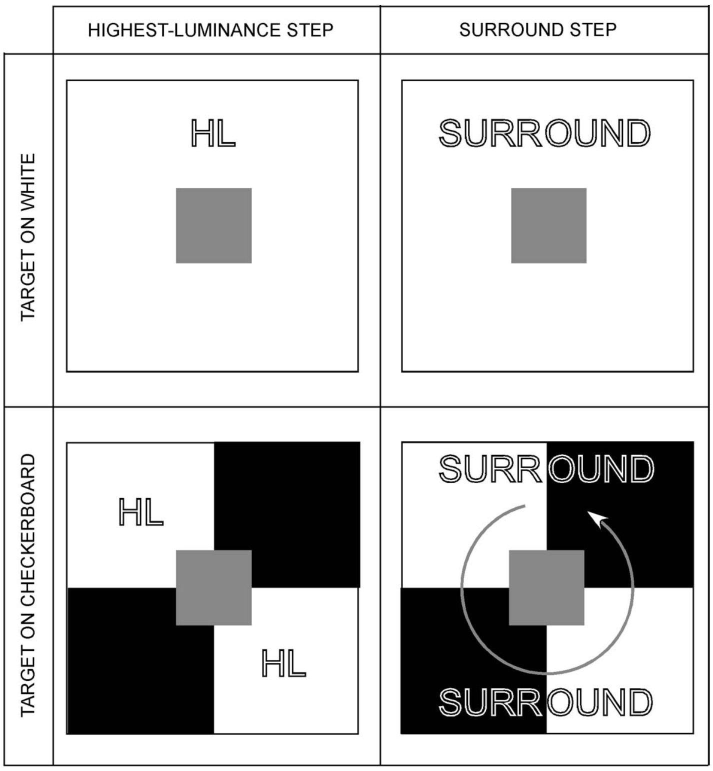

Figure 6. Explanation of why a target on a white background darkens relative to its luminance match set on a black-and-white checkerboard. The diagram indicates which regions serve as highest luminances (left column) and surrounds (right column) for the two targets. The highest luminance is the same for the two targets, but the surround is not: It is the white background for the target on white, but the average luminance of the black and white checks for the target on the checkerboard. At the surround step, the target on white has a lower luminance ratio to the surround and thus receives a lower lightness assignment.

A grey target sitting on a white background is computed as grey at both local steps (Figure 6, top row); but if it sits on a black and white checkerboard, the assignments at the two steps are going to be different (Figure 6, bottom row). At the HL step, the target is assessed relative to the white checks in the checkerboard, and assigned a value of grey. At the surround step it is assessed relative to the black and white checks, or, equivalently (since they cover equal areas), to their average luminance, which is necessarily lower than the luminance of white. Thus, the target on the checkerboard is identical to the target on white at the HL step, but lighter at the surround step. Accordingly, the target on white is expected to look slightly darker than its "true" match, as observed experimentally, and there is no need to assume a scale normalization factor.

Simultaneous Contrast with Articulated Surrounds

Simultaneous lightness contrast is stronger when the dark and light backgrounds of the classic display of Figure 1a are replaced by articulated fields, as in Figure 7. Each articulated field is a checkerboard obtained by breaking the uniform surround into many different square regions, some brighter and some dimmer, so that there is no change in the average luminance (see also Adelson, 2000). Enhanced contrast with articulated surrounds has been shown to be the outcome of two separate events: (1) a lightening of the incremental target, and (2) a darkening of the decremental target (Bressan & Actis-Grosso, in press; see also Arend & Goldstein, 1987; Schirillo, 1999). Whereas the latter effect necessarily results from the change in highest luminance that follows surround articulation, the former does not, indicating that surround articulation as such plays a role in lightness assessment.

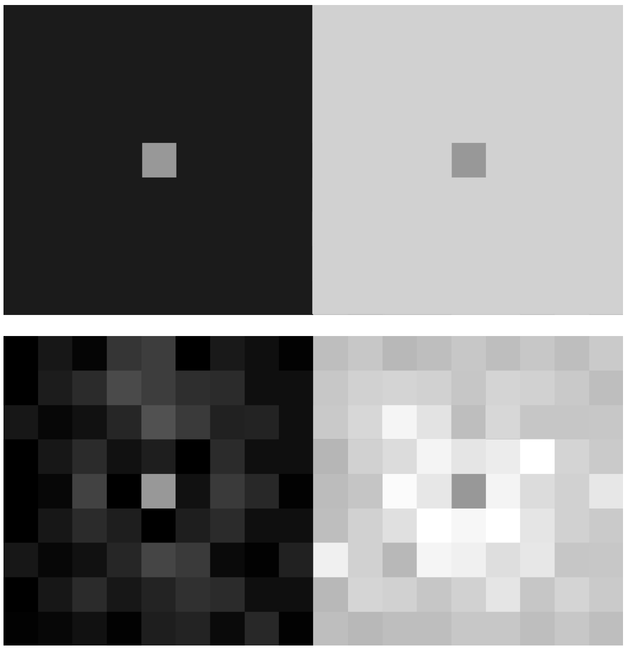

Figure 7. Simultaneous lightness contrast is stronger with an articulated surround. The decremental target on the articulated surround (bottom right) darkens relative to the same target on the plain surround (top right), because its local framework contains a highest luminance. The incremental target on the articulated surround (bottom left) lightens relative to the same target on the plain surround (top left), because it receives a higher superwhite assignment at the surround step.

The original anchoring model proposes that articulation augments simultaneous lightness contrast by increasing the weight of local anchoring relative to global anchoring. But this cannot be true. The target on the articulated surround lightens, relative to an identical target on the uniform surround, even when it is the highest luminance in the visual field (Bressan & Actis-Grosso, in press). Being at the same time local and global anchor, each target must receive identical values locally and globally. Thus, altering the local weight relative to the global cannot possibly make a difference on the final weighted average.

From a double-anchoring stance, surround articulation behaves like surround size or surround luminance, and affects the weight assigned to the surround step relative to the HL step (*weight/articulation* rule). Incremental patches are white at the HL step and superwhite at the surround step: articulation increases the weight of the local superwhite relative to the local white. This results in a lightening of the incremental target on the articulated surround relative to the same target on the plain surround, as observed.

An important implication of the *weight/articulation* rule is that, via within-framework weighting only (that is, if the highest luminance does not vary), changes in surround articulation are expected to affect full or partial increments, but not full decrements. The reason is that decrements receive the same assignment at both steps, so that shifts in surround-step weight relative to HL-step weight are inconsequential. This prediction is consistent with the otherwise unexplained observation that, unlike matches to increments, matches to decrements are virtually unaffected by the heterogeneity of the checkerboard (Schirillo & Shevell, 1996).

Incidentally, the significant lightening of the incremental target on the articulated surround independently confirms the inability of the original anchoring theory to account for double-increment illusions. In the original model, the lightening of a highest luminance can be explained only by invoking either the area rule or a scale normalization factor, but in this case both predict an effect in the wrong direction. According to the area rule, the larger it is, the more a surround tends to white and pushes the target towards luminosity. However, the areas of the plain and articulated surrounds are equal; if anything, the larger uniform area is the plain surround. According to scale normalization, the perceived difference between the dark background and the light target increases, and this expansion is expressed not only as a further darkening of the background, but also as a further lightening of the target. Yet this implies more lightening in the framework where the range of luminances is smaller, on the uniform surround that is, which is again the opposite of what is observed.

Remote Contrast

The lightness of a figure can be affected by regions that do not share any borders with it. The most venerable illustration of what we may name *remote contrast* dates back to Wolff (1934). His display (a variant of which was more recently studied by Reid and Shapley (1988), who called the effect "assimilation") consists of two squares of equal luminance, surrounded by square frames that also have equal luminance, surrounded in turn by square frames of different luminance. The square whose outer frame is more luminant looks darker.

I offer an especially challenging variant of remote contrast in Figure 8. Here, two identical grey disks are surrounded by identical white annuli, placed on a luminance gradient. The disk whose (distant) surround is represented by the dark end of the gradient looks lighter than the other. The original anchoring theory cannot handle this illusion, because the lightnesses of the two targets are assessed relative to identical anchors (the white annuli), and are thus predicted to be equal.

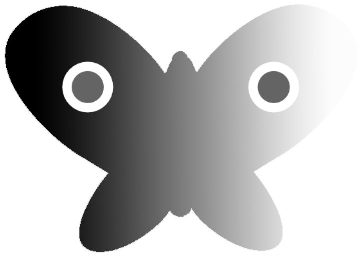

Figure 8. The butterfly illusion: remote contrast with identical white annuli. The disks are the same shade of gray, but the one on the left seems slightly lighter.

For each grey disk in Figure 8, one can identify a local framework based on adjacency (the white annulus), and a super-local framework based on proximity. For the disk on the left, this framework consists in the dark half of the gradient; for the one on the right, in the light half of the gradient. Relative to the disk on the right, the disk on the left receives a higher lightness value at both steps, within its super-local framework; hence, it will look slightly lighter, as observed.

Reverse Contrast or "Assimilation"

There are displays where simultaneous contrast is overturned completely and replaced by an opposite effect, whereby a target surrounded by more white than black looks lighter rather than darker. Reverse contrast effects arise naturally within a framework approach: ordinary contrast between the target and a remote group of elements with which the target groups strongly can, of course, overcome ordinary contrast between the target and a local surround with which the target groups only weakly. Here we will deal with two examples, Bressan's dungeon illusion (2001) and bullseye displays (e.g. Bindman & Chubb, 2004).

The dungeon illusion. The display in Figure 9, top panel, is a stronger variant of Bressan's (2001) dungeon illusion. (The original version used squares rather than disks, and the space between them resembled the grid of a dungeon's window, partially occluding the grey target regions underneath.) The grey disks on the white background look lighter than the identical grey disks on the black background, an incomprehensible outcome from a strictly local contrast perspective. From a double-anchoring stance, each set of grey disks participates in two separate frameworks, one including the surrounding disks and one including the background. The first framework is founded on luminance polarity (all disks have the same contrast sign relative to their common surround) and shape similarity. The second framework is founded on retinal proximity only. (Retinal proximity is the weakest among grouping factors: it merely means simultaneous presence in the visual field.) Hence the target disks are expected to be influenced by both the contextual disks and the background, but by the contextual disks more strongly. In this framework, the target disks on the left are white at the HL step and superwhite at the surround step; the target disks on the right are grey at both steps.

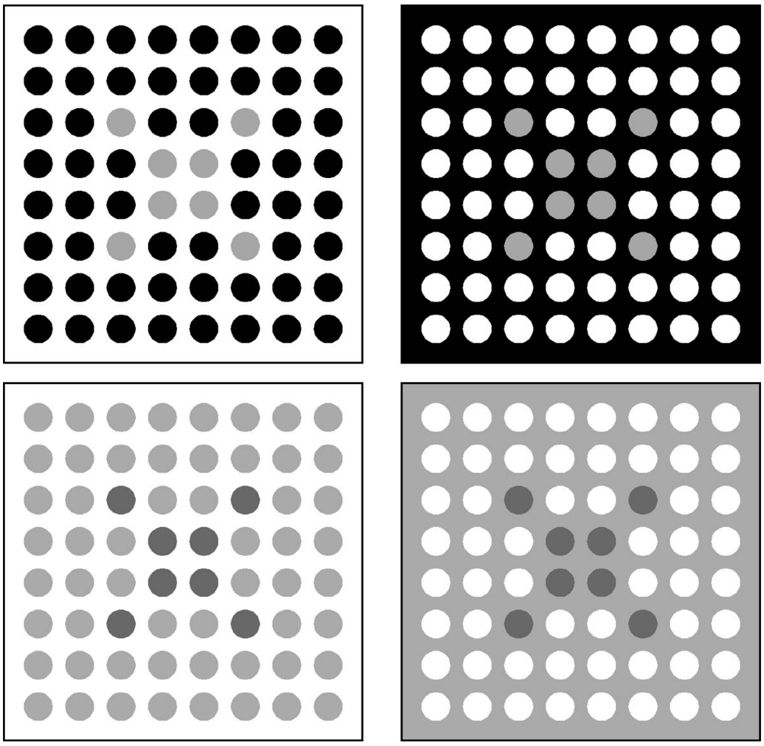

Figure 9. Top panel: Reverse lightness contrast in a version of Bressan's (2001) dungeon illusion. The central gray disks in the left display are entirely surrounded by white, but they look lighter than the identical central gray disks in the right display that are entirely surrounded by black. Bottom panel: Although the spatial structure of the displays is unchanged, reverse contrast is replaced by ordinary contrast. The central disks in the left display, entirely surrounded by white, look darker than the identical central gray disks in the right display, entirely surrounded by gray.

The double-decrement variant in the bottom half of Figure 9 lends separate support to this argument. Stronger grouping with the contextual disks would produce here reverse contrast as well, by darkening the target disks more on the right (where the contextual disks are white) than on the left (where they are light grey). Yet, it is the target disks on the left that appear darker: reverse contrast is abolished, and replaced by ordinary contrast. In the model, this follows from the dramatic weakening of the contextual-disk framework on the right, where the grouping factor of luminance polarity has been destroyed (target disks and contextual disks have opposite contrast signs relative to their common grey surround), and luminance similarity actually encourages grouping with the background.

As it stands, the original anchoring model (that does not include luminance polarity and similarity among its grouping factors) is unable to predict these two asymmetrical results. In the four displays of Figure 9, the highest luminances are identical; the areas are identical; the luminance ranges are identical; thus, however frameworks are chosen, contrast should go in the same direction (i.e., be either ordinary or reverse) in both the top and bottom pairs.

Bullseye displays. A grey patch surrounded by rings that alternate outward from black to white looks darker than an identical patch surrounded by rings that alternate outward from white to black, as shown in Figure 10, top panel. It has been claimed that such an effect stems from "assimilation" rather than contrast, and cannot possibly be explained by anchoring theories. "Assimilation" would occur for a given region whenever the contrast at its edge (here, at the border between the grey disk and the innermost ring) is small compared to the contrast of edges in the general neighborhood (here, at the borders between black and white rings), and the density of edges in the neighborhood is high (Bindman & Chubb, 2004).

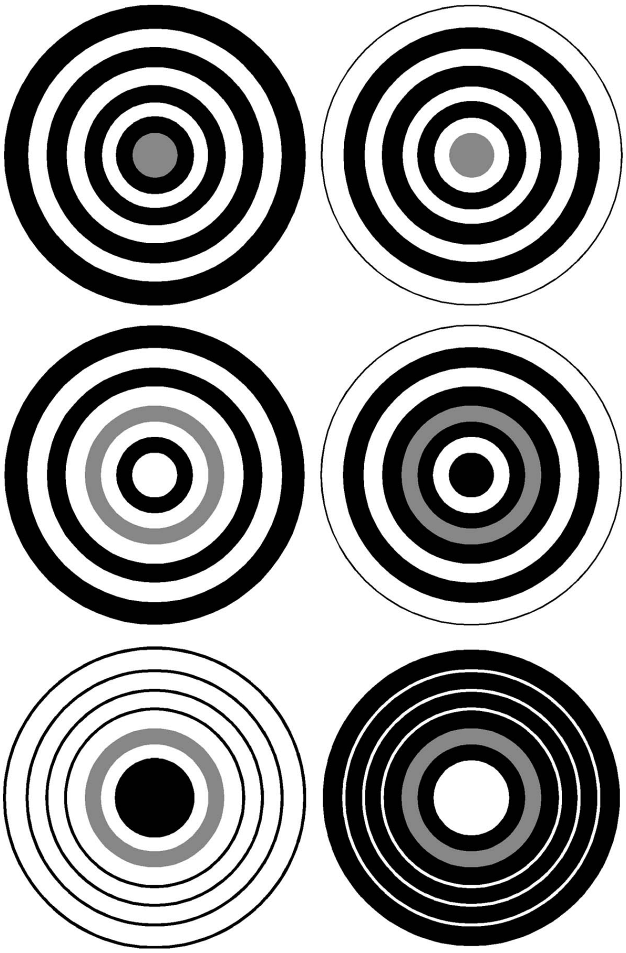

Figure 10. Reverse lightness contrast with bullseye displays. The top display is similar to one studied by Oh and Kim (2004). The gray disk on the left (surrounded by black) looks paradoxically darker than the identical gray disk on the right (surrounded by white). Middle display: Turning the central disks into rings, and thereby increasing grouping by "alternation," enhances the effect. The gray ring on the right (surrounded by black) looks darker than the identical gray ring on the left (surrounded by white). Bottom display: The ring surrounded by black does not look darker than the ring surrounded by white, although the gray rings have the same local contrast relationships as the corresponding gray rings in the middle panel.

However, given the resemblance between this effect and the dungeon illusion (with rings in the role of contextual disks), invoking "assimilation" seems unnecessary. On the basis of independent data, it has in fact been suggested (Oh & Kim, 2004) that, by alternation, the central disk groups with the set of white rings in the left display, and with the set of black rings in the right display. Yet a disk is not a ring. If the illusion is based on grouping with the rings, introducing the factor of shape similarity should strengthen it. I have done this in the middle panel of Figure 10. Clearly, a more obvious grouping produces a more striking effect.[5]

In the bottom panel, I have eliminated the rings framework while preserving the conditions for "assimilation", that is the local contrast relationships for each disk and the high contrast and density of edges in the neighborhood. For most observers, the "assimilation" effect is gone.

With this kind of reverse-contrast displays, it has been a consistent and surprising finding that informal presentation or a forced-choice procedure induce compelling "assimilation" effects, but these disappear (and are sometimes replaced by regular contrast) with an adjustment procedure (Bindman & Chubb, 2004; De Weert & Spillmann, 1995; Oh & Kim, 2004). This is only to be expected if reverse contrast is founded on remote grouping. When people are asked to adjust a separate patch till it matches the achromatic color of the target region, their attentional window only includes the target and, at most, some of its immediate surround. This is bound to produce, respectively, either no effect or a regular contrast effect, as observed experimentally.

Testing the Model: Figure vs Ground

The lightness of a region can change depending solely on whether such a region plays the role of figure or ground. Figure 11 offers a variant of a display that appears in Wolff (1934). The disks on either side of the display have the same luminance as the background on the other side; but the dark disks look very slightly darker than the dark background, and the light disks look very slightly lighter than the light background. The effect is not contingent on the different sizes and shapes of figures as opposed to backgrounds, as shown by the fact that, in a reversible display, the same area varies in lightness as a function of whether it is interpreted as figure or ground (Coren, 1969).

As Figure 11 shows, this lightness modulation is easy to explain within a double-anchoring approach. Locally, the dark disks receive a *lower* assignment than the physically identical dark background (because the latter is white at the surround step). Locally, the light disks receive a *higher* assignment than the physically identical light background (because the former are superwhite at the surround step). Both effects arise at the surround step. Since adding a superwhite to a white entails a smaller difference than adding a white to a dark grey, the model also predicts the further empirical subtlety that the effect is stronger for the dark regions than for the light ones.

[5] A display similar to those in Figure 10, middle panel, was used by Howe (2005) to show that White's effect survives removal of T-junctions. Although his point is well taken, the problem with Howe's display is that its rings were much finer and actually met the conditions for assimilation (in terms of neuronal summation or averaging), as shown by the fact that the direction of the illusion was the same even when all contextual rings were removed.

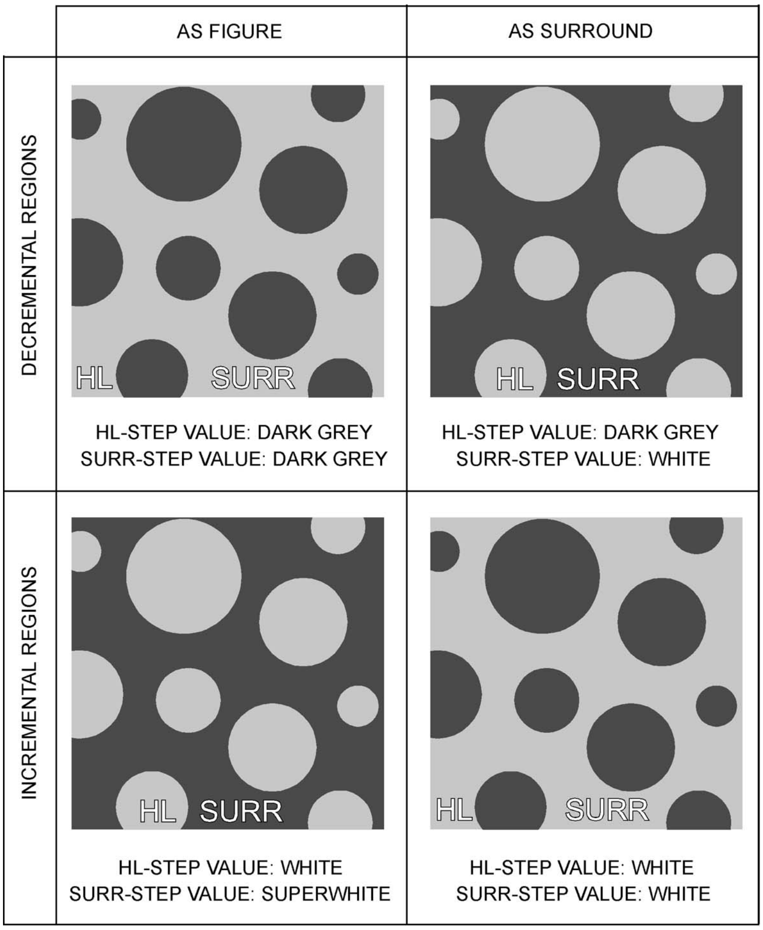

Figure 11. Double-anchoring explanation of the Wolff illusion. Local values are shown for each decremental and incremental region, at each step, as a function of whether such a region is seen as a figure or a background. HL = highest luminance; SURR = surround.

All the action takes place in the local frameworks. In more peripheral frameworks, assignments are in practice identical for all the light regions and identical for all the dark regions, and will water down the illusion considerably; accordingly, the lightness differences in Figure 11 are weak at best. But if each local framework completely filled the visual field, the peripheral frameworks would not be there any longer, and the illusion would be striking.

Testing the Model: The Gelb Effect and its Variants

A physically black disk, suspended in midair and illuminated by a beam of light in a half-darkened room, looks white (Gelb, 1929) or even luminous. When a larger surface of higher luminance, such as a piece of real white paper, is brought into the beam and placed behind the disk so as to surround it, the latter becomes perceptually black. This only happens when the white paper completely surrounds the disk: if it is simply held in front of it, the black disk appears middle to light grey. Perceived reflectance values ranging from 18.9 to 24.6 (Newson, 1958), from 41.6 to 68.4 (Stewart, 1959), and from 33.7 to 52.3 (Gilchrist et al., 1999) have been reported.

The double-anchoring model predicts that the Gelb disk, on its own, will look white, superwhite, or luminous. The target is always white at the HL step, and superwhite relative to the dark surrounding void; the importance of the superwhite component in the final computation, and thus the disk's actual appearance, will depend on the surround-step weight, and therefore on (a) the grouping factors that link the target to its surround (only retinal proximity if the disk is suspended in midair), and (b) the relative area, articulation, and absolute luminance of the surround.

Likewise, the double-anchoring model predicts that the target will appear black when completely surrounded by white. In this case, the values computed at the surround and HL steps are identical, and weighting inconsequential; no variability is expected. When the white paper does not surround the target, but is suspended in front of it (as in Stewart, 1959), on the other hand, the model predicts that the black paper will look grey, as a compromise between superwhite (its surround-step value) and black (its HL-step value). The exact shade of grey will depend on how the two steps are weighted, hence on the factors listed in the paragraph above; the resulting variability is consistent with the large range of perceived greys obtained, with this kind of display, by different investigators in different conditions.

Other conditions being equal, the relative weight of the HL step will depend on the size of the white paper (*weight/area* rule), a prediction that has received empirical support (Stewart, 1959). It will also depend on the proximity between the white paper and the target, which agrees with empirical data showing that the effect of a highest-luminance region on the appearance of a test region depends on how close they are (Newson, 1958). If this distance is zero (the two regions touch), the highest-luminance region becomes part of the target's local surround, thereby changing the target/surround ratio used at the surround step, and additionally darkening the target. This is consistent with the finding that, when the two regions are adjacent, moving the highest-luminance surface away through the first millimeter produces a greater effect than moving it through the next 60 cm (Newson, 1958).

Staircase Gelb Effect

Imagine that the black disk in a spotlight of the Gelb effect is replaced by a strip of five contiguous squares, going from black to white in approximately equal steps. The perceived lightness range for the five squares is compressed relative to the actual range, as can be seen in the left panel of Figure 12 (Cataliotti, 1993; Gilchrist et al., 1999). The original anchoring model accounts for this result by assuming a compromise between the “veridical” data predicted by highest-luminance anchoring (the diagonal line in the left panel of Figure 12) and the fact that each of the grey squares by itself would be the highest luminance in the global framework, and thus appear white (the horizontal line in the left panel of Figure 12). However, none of the grey squares is actually the highest luminance in the global framework, because the latter contains the white square also, and it is the white square that is the highest luminance. Therefore, it is not clear why this line should be drawn as horizontal. The model postulates that this happens because the white square is small relative to the entire laboratory scene, and is therefore a weak anchor. Note, however, that the global framework line would be horizontal only if the effect of the white square was actually zero.

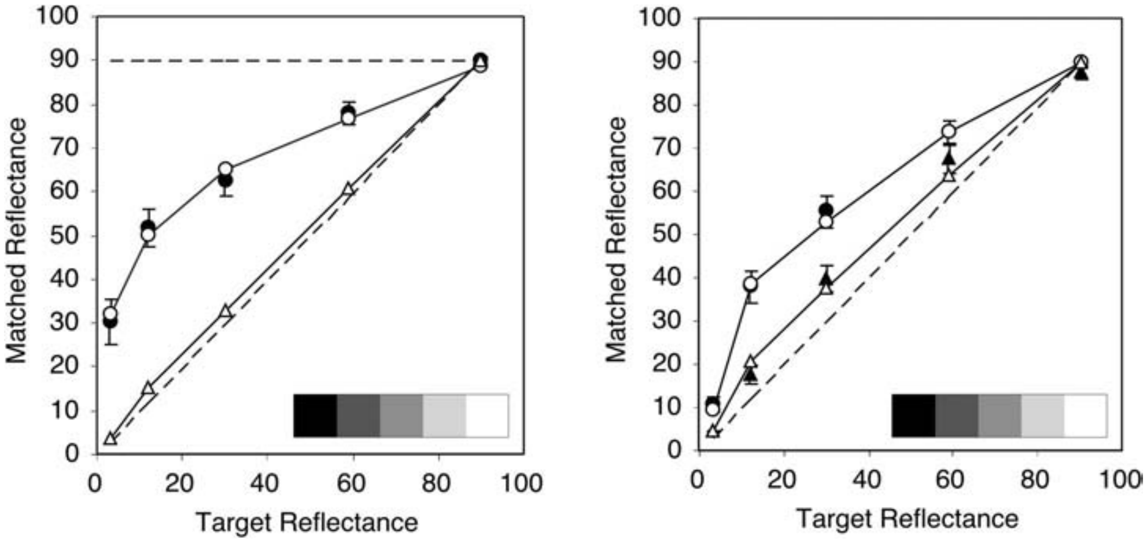

Figure 12. Left panel: The staircase Gelb effect as observed (solid symbols) and as predicted by the double-anchoring model (open symbols) in a multiluminance laboratory. Circles indicate spotlight on, and triangles indicate spotlight off (Cataliotti, 1993). The dashed diagonal line shows the "veridical" matches. When the spotlight is removed, a fivefold decrease in surround-step weight is enough to make compression virtually disappear (as can be seen by the coincidence between the open-triangle line and the diagonal line), because it is combined with an increase in the weight of the peripheral framework, where anchoring is "veridical." Right panel: The staircase Gelb effect as observed (solid symbols) and as predicted by the double-anchoring model (open symbols) in a black laboratory. Circles indicate spotlight on, and triangles indicate spotlight off (Bressan, 2002.) The dashed diagonal line shows the "veridical" matches. When the spotlight is removed, a fivefold decrease in surround-step weight reduces, but does not eliminate, compression, because there is no "veridical" anchoring in the peripheral framework. In the double-anchoring model, the lightness corresponding to the square of 90% reflectance in the spotlight always results in superwhite, which agrees with the data reported by Cataliotti and Gilchrist (1995). For convenience, values larger than white are represented as white in both panels of Figure 12. See Appendix B.

The account of the staircase Gelb effect given by the double-anchoring model is quite different. In the original model, compression is due to global, as opposed to local, anchoring—whereas local anchoring is assumed to be "veridical". In the double-anchoring model, the basic effect is actually provided by local anchoring, and arises at the surround step. Say that each square belongs to two frameworks: a powerful local one (the row of squares in the spotlight), by way of the grouping factors of luminance polarity and similarity, coplanarity, proximity, and alignment; and a weak peripheral one (the rest of the visual scene), by way of retinal proximity only.

If the laboratory is less illuminated than the squares, in the peripheral framework all squares are white at the HL step, so the peripheral component does contribute to their apparent lightening. Because of the unbalance in grouping strengths, however, in the squares-under-spotlight condition such a framework is expected to carry little weight.

What plays local surround for each square in the local framework? Each square is joined to the two adjacent squares via the strong variants of alignment and proximity, i.e. T-junctions and adjacency. Hence, the local surround of each square is equivalent, in practice, to the average luminance of the two adjacent squares, the one on its left and the one on its right. We may well assume that the strength with which a surface groups with non-adjacent regions—and then, their weight in the final computation—is an inverse function of the square of their distance, or some similar measure (see Zaidi, Yoshimi, Flanigan, & Canova, 1992); given the paucity of their contribution, then, non-adjacent squares will not be entered in the equation. As to the first and last squares, they group with a single adjacent square rather than two: for them, therefore, we shall reduce the weight of the surround step by half.

The solid symbols in the right panel of Figure 12 show the perceived reflectances of five linearly arranged squares of reflectances 3, 12, 30, 59, and 90%, suspended in midair in a laboratory entirely painted flat black, in two separate conditions (Bressan, 2002; Dal Toè, 2002). The top curve, solid circles, illustrates the standard effect, obtained by illuminating the five squares with a 650-watt spotlight, and having ten observers match each square to one of the chips in a 16-step Munsell scale, placed inside a box and separately illuminated. The open circles represent the values generated by Equation 2, assuming equal weights for the surround and HL steps.

To show that the bulk of the effect obtains in the local framework, the weight given here to the peripheral framework is zero; larger weights, other conditions equal, would displace the curve upwards, further increasing compression. Note that the observed curve is similar to the one reported by Gilchrist et al. (1999) and represented in the left panel, although their laboratory contained a 30:1 range of luminances, with the highest luminance equal to the luminance of the 3% square in the spotlight. The large absolute amount of compression and the similarity between the two curves agree with the idea that the peripheral framework plays here a supporting role only. Such a supporting role is well visible, however, and is illustrated by the fact that the points in the top curve (right panel) are lower than the corresponding points in Gilchrist et al.'s curve (left panel). This happens because the peripheral framework weighs considerably more in Gilchrist et al.'s experiments than in mine, since their laboratory was much larger (28 m^2 vs 9 m^2), much more illuminated (room lights on vs off), and especially much more articulated (full range of reflectances vs one reflectance). Likewise, Gilchrist et al. (see Cataliotti, 1993) did find a slight decrease in compression, that is lower lightness matches, when their staircase experiment was replicated either with room illumination off, or with the laboratory masked off from view by a booth. This puzzling failure to get rid of the effect to any substantial extent was attributed to the impossibility of eliminating the outer framework entirely, but such a failure makes complete sense if one considers that the outer framework contributes to, but does not totally determine, compression.

When the special lighting on the squares is removed, in Gilchrist et al.'s experiments nearly all of the compression disappears as well (A. Gilchrist, personal communication, June 2000; quantitative data are not available). The spotlight-off condition differs from the spotlight-on condition in two respects. The first is a sharp decrease in the absolute luminance level on the squares, and therefore in the luminance of each square's surround. In keeping with the *weight/luminance* rule, we shall then expect a smaller weighting at the surround step. (For independent evidence for the importance of absolute luminance in lightess computation, see

Masin, 2003b.) The second is that the luminance range of the squares and the luminance range in the laboratory become the same. This has two consequences. First, the powerful grouping factors of luminance polarity and similarity that kept the squares in the spotlight together are removed, and the local: peripheral balance reverses. Second, in the peripheral framework the highest-luminance anchor is played not by each square, as before, but by the white regions in the laboratory. Therefore, peripheral anchoring becomes "veridical". In the left panel of Figure 12, the open triangles represent the values generated by Equation 2, under a five-fold decrease of surround-step weight relative to HL-step weight, and a five-fold increase of peripheral-framework weight relative to local-framework weight. Since there are no quantitative data to model, the choice of these weights is arbitrary, but the important point is that diminishing the contribution of surround anchoring and augmenting that of peripheral anchoring both have the same effect, that of cutting down compression.

In my experiments, on the other hand, the luminance of surfaces in the black laboratory is still lower than that of the squares when the spotlight is switched off (and normal room illumination is switched on). The only difference between the spotlight-on and spotlight-off conditions is a drop in the luminance of each square's surround. The bottom curve in the right panel of Figure 12, solid triangles, depicts the matches produced, by a different group of ten observers, under normal room illumination (same level of illumination falling on the squares and on the Munsell scale). As can be seen, a considerable amount of compression survives the removal of the spotlight. The open triangles represent the values generated by Equation 2, under a five-fold decrease of surround-step weight relative to HL-step weight.

The original model cannot account for these data. Since my laboratory is completely black, the lightness of each of the squares should still be a compromise between its "veridical" value in the local framework and its white value in the global framework, thus the before- and after-spotlight curves should coincide. The explanation could only be rescued by assuming a shift in weight from the global to the local framework, but the global framework's articulation does not change when the spotlight-in-the dark is replaced by normal room illumination, and, if it changes at all, it most certainly cannot diminish. Switching off the spotlight on the squares should, if anything, decrease their segregation from the rest of the laboratory, making the local framework weaker, not stronger.

Hypercompression

When the spotlight on the five squares is made brighter, compression becomes larger (A. Gilchrist, personal communication, June 2000). From the standpoint of the original anchoring model, this should not occur. Once the spotlight is bright enough to produce a horizontal line for the global values, compression is as large as it can get, because in the model the compression comes from a compromise between the local line (diagonal in Figure 12) and the global line (horizontal in Figure 12). The only other possibility is a shift in weight from the local to the global framework, but there is no reason for this. If anything, a brighter spotlight should segregate the local framework further, making it stronger.

In the double-anchoring model, hypercompression is a direct consequence of the shift in weight between the two anchoring steps, following an increase in surround luminance level (*weight/luminance* rule). Augmenting the relative weight of surround anchoring makes the squares lighter, increasing compression.

Incidentally, the *weight/luminance* rule implies that abrupt changes in surround luminance (when ratios do not change) are expected to affect increments, but not decrements. Figural decrements receive the same assignment at both steps, so relative weighting is unimportant. This is consistent with the empirical observation that, as illumination varies, decrements remain relatively constant, while increments do not (Arend & Goldstein, 1987).

Insulation

Gilchrist and Cataliotti (1994) found that a white border, placed under the spotlight so as to surround the row of squares of the staircase Gelb effect, seems to insulate the squares from the influence of the global framework. This is how these authors describe the fact that all the squares appeared darker when surrounded by white, as opposed to black (or to darkness), as shown in the left panel of Figure 13. The effect receives no explanation within the original anchoring model, because there is no change in the highest luminance.

From a double-anchoring standpoint, the explanation is simple: the white border represents a new, powerful framework, to which each of the squares belongs via adjacency, depth similarity, luminance polarity and similarity. Let's consider five linearly arranged squares of reflectances 3, 12, 30, 59, and 90% suspended in a void. In the right panel of Figure 13, the top curve illustrates a standard staircase Gelb effect, as generated by Equation 2. The bottom curve illustrates the same case, with the difference that the dark peripheral framework, where all squares are the highest luminance, has been replaced by a white peripheral framework, where anchoring is “veridical”. A comparison between the two curves shows that, in order to explain the effect, the concept of insulation is unnecessary.

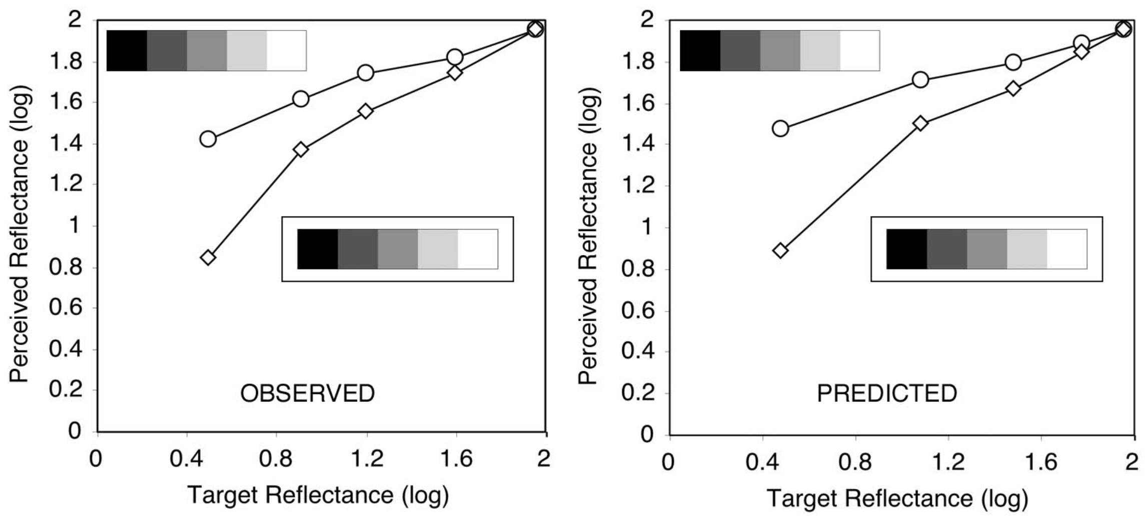

Figure 13. Left panel: Insulation. The two curves represent the lightness values observed in a no-border (upper line) and white-border (lower line) staircase Gelb effect (Gilchrist et al., 1999). Right panel: The two curves represent the lightness values predicted by the double-anchoring model in a no-border staircase Gelb effect (as in Appendix B, top block) and in the corresponding white-border case.

Testing the Model: The Perception of Luminosity

In the double-anchoring model, a figure whose final lightness value is larger than 90 (perceived reflectance of white) looks *superwhite*, that is either as a glowing or especially illuminated white, or luminous. This condition is not meant to cover all the instances in which an object appears to emit light. In many cases, the impression of luminosity is induced by the presence of a surrounding gradient (see Zavagno, 1999), and entails no major lightness change. An example is the double-diamond illusion presented in Bressan (2001), where a white region surrounded by the dark end of a gradient appears lighter than an equal region surrounded by the light end of a gradient, but it is the latter that seems to emit diffused light. Another example is Figure 2: although the tops of the blocks that represent a smaller increment relative to the surrounding gradient may look luminescent, on the lightness scale they occupy a lower place than the tops representing a larger increment. In this paper we are concerned not with luminosity as the impression of a light source, but with luminosity as superwhite. Note however that, in a natural achromatic world, most light sources do entail superwhite—insofar as they are usually large luminance increments relative to the rest of the scene.

Effect of Size on the Luminosity Threshold

The threshold luminance value at which an incremental target begins to appear luminous increases with its size relative to the surround (Bonato & Gilchrist, 1999). In other words, if we keep the area of the surround constant, a small target is seen as luminous more easily than a large target; which amounts to stating that a target sitting on a large surround is seen as luminous more easily than an identical target on a small surround. Within the double-anchoring model, this result is readily explained: targets on large surrounds appear lighter than targets on small surrounds, because larger surrounds weigh more at the surround step. Since incremental targets receive a white assignment at the HL step and a superwhite assignment at the surround step, targets on large surrounds will approach the luminosity threshold earlier than targets on small surrounds.

A parallel result is that the (photometrically identical) surrounds on which these targets lie are seen as having different lightnesses. Bonato and Gilchrist (1999) showed that a white surround was seen as off-white regardless of the actual luminance of a small incremental square sitting on it (solid triangles in Figure 14). However, the lightness of the same surround was an inverse function of the square's luminance when the incremental target square was much larger (solid circles in Figure 14). The open symbols show the predictions of the double-anchoring model for the small-target and large-target conditions, when the luminances used in the formula are exactly those of Bonato and Gilchrist's stimuli, and the weight given to the surround step is proportional to its relative area (1 for the large-target condition, 17 for the small-target condition).

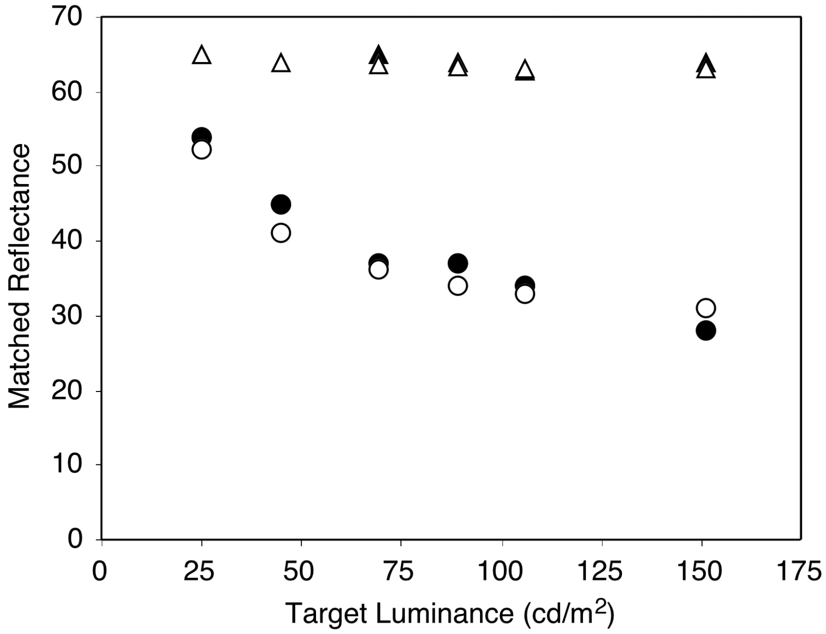

Figure 14. The lightness of a background (121 deg^2 of visual angle in area) is much more affected by the luminance of an incremental square sitting on it when this square is large (68 deg^2 in area; solid circles) than when it is small (4 deg^2 in area; solid triangles). Adapted from "Perceived Area and the Luminosity Threshold," by F. Bonato and A. L. Gilchrist, 1999, *Perception & Psychophysics, 61,* p. 790. The open symbols represent the values predicted by the double-anchoring model (see Appendix C).

Effect of Figure vs Ground on the Luminosity Threshold

A region perceived as figure has a lower luminosity threshold than a region perceived as ground, even when the areas of the two regions are equal (Bonato & Cataliotti, 2000). For example, a light region seen as a profile against a black background begins to appear luminous at 29 cd/m^2, but the same region seen as the background of a black profile begins to appear luminous only at 67 cd/m^2. The original anchoring model accounts for this result by suggesting that these areas are not really equal: the area of the ground is functionally greater, because it is perceived to extend behind the figure. Thus we may expect that, by virtue of the area rule, the lightness of the background goes up, and at the same time pushes the incremental figure towards superwhite.

In the double-anchoring model, there is no need to presume that surrounds must have a different area than measured. (As we shall see in the next section, sometimes this assumption leads to the wrong prediction also). An incremental region seen as surround yields a local value of white (it is white at the HL step, and white at the surround step), but the same region seen as figure yields a local value of superwhite (it is white at the HL step, and superwhite at the surround step). The lower the luminance of the surrounding region, the larger this superwhite assignment is, determining a lower luminosity threshold.

Testing the Model: Lightness in a Dome

If the lightness of every single region in our visual world depends on its relationship to all the others, one feels that even the simplest laboratory conditions are too cluttered to begin deriving our models' basic rules. Li and Gilchrist (1999) took a clever step away from such complications by placing their observers' heads in a large illuminated hemispherical dome, the inside of which was divided into two parts, painted black and middle grey. An arrangement such as this meets the minimal conditions for perceiving a surface (Koffka, 1935, p. 111): two regions that fill the entire visual field, separated by a single border. Schematic representations of what observers saw inside the dome are shown in Figure 15.

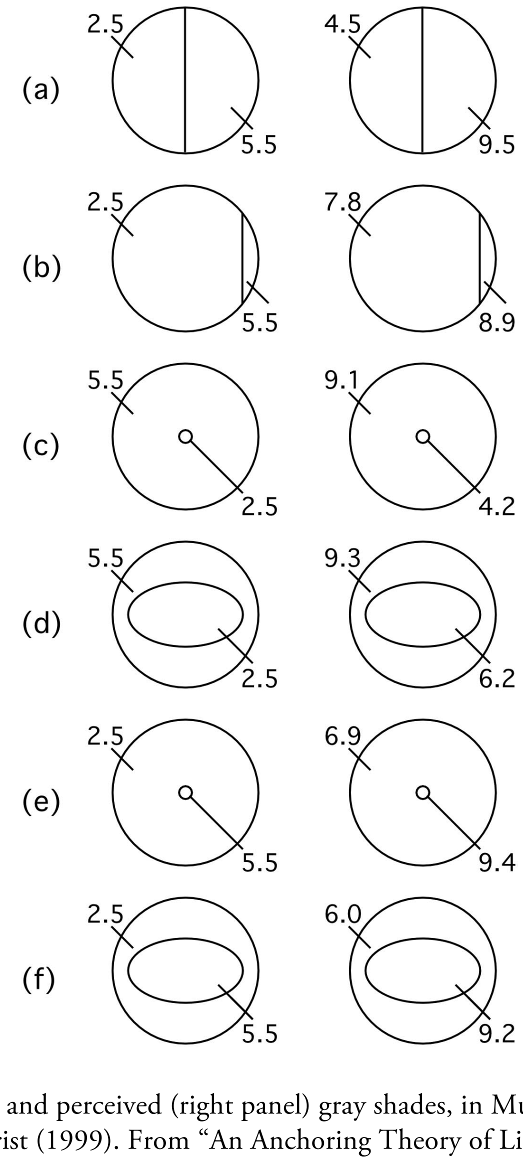

Figure 15. Actual (left panel) and perceived (right panel) gray shades, in Munsell values, in the dome experiments of Li and Gilchrist (1999). From "An Anchoring Theory of Lightness Perception" by A. Gilchrist, C. Kossyfidis, F. Bonato, T. Agostini, J. Cataliotti, X. Li, et al., 1999, *Psychological Review, 106,* p. 800. Copyright 1999 by the American Psychological Association. Adapted with permission.

Results Apparently Undermining the Surround-as-White Rule

Here, we will consider the results that Li and Gilchrist (1999) interpreted as undermining the surround-as-white rule, and show that they are instead consistent with such a rule, when this is assisted by the highest-luminance principle—as it is in the double-anchoring model. These results are: backgrounds do not always appear white; backgrounds that surround a figure do not appear lighter than backgrounds that are simply adjacent to a figure; a surface seen as background does not appear lighter than a photometrically identical surface (allegedly) seen as figure.

Backgrounds do not always appear white. In the incremental large oval condition (Figure 15f), the black background appeared middle grey (Munsell 6) and not white. This is not a problem for the double-anchoring model: surrounds do not necessarily appear white, because they are also anchored to the highest luminance in the scene. That is, whereas according to the surround rule as proposed by Gilchrist and Bonato (1995) the surround *must* appear white (at least in the local framework), in the double-anchoring model it need not appear white—it only must function as white in the computation.

Surrounding backgrounds do not appear lighter than adjacent backgrounds. Compare the black region in the unevenly split dome (Figure 15b) with the black region in the disk/surround display (Figure 15e). Li and Gilchrist argue that, if there were a lightening effect in a surround in addition to that of its area, the black region would look lighter in the second condition, where it completely surrounds the disk: but the contrary is observed (Munsell values 7.8 vs 6.9, a significant difference).

In the double-anchoring model, the surround-as-white rule implies that all surrounds are given a default value of white, not that enclosing surrounds are given a lighter value than non-enclosing surrounds. But how can one explain the direction of the difference? Note that the surround in Figure 15b almost completely fills the observer's visual field (the whole left semi-field plus 59 degrees on the right), and the highest luminance is confined to a very peripheral area, since the observer was asked to fixate straight ahead. As long as this extremely peripheral highest luminance counts less than the foveal highest luminance in the disk/surround dome of Figure 15e, the model does indeed predict that the black area looks lighter in Figure 15b than in Figure 15e.

On the basis of the highest-luminance rule, the anchoring model would predict equal values for the two black regions. Since the black area is, if anything, greater in the disk/surround dome, engaging the area rule would lead to the prediction that the black region appears lighter in this condition, the opposite of what is observed.

Background does not appear lighter than (alleged) figure. The last apparent piece of evidence against the surround rule comes from the comparison between the black regions in the two large oval conditions. In Figure 15f, the black region represents the background, and as such it should appear, if not white, at least lighter than the oval black region in Figure 15d. On the contrary, they look identical (Munsell 6 vs 6.2, a non-significant difference).

Now, it is highly questionable whether the oval in Figure 15d can be regarded as a perceptual figure on a background. In these experiments, the oval area subtended 118 x 91 degrees of visual angle, and observers were asked to fixate in the center. That is, this was a huge region rimmed by a relatively thin (and very peripheral) border: much more like a large, empty surround than a figure. Incidentally, if it had been a perceptual figure, the oval in Figure 15d would have appeared significantly darker than the background in Figure 15f, because of the well-known Wolff effect (illustrated earlier in this article). If this is the case, then, in a double-anchoring approach the black regions of Figures 15f and 15d are indeed expected to appear identical.

Li and Gilchrist interpret the whole of their results as supporting the area rule (the larger, the lighter), rather than the surround rule. Note however that, if the area rule made sense, the black region in Figure 15f should look the same shade of grey as the black region in Figure 15a, since they are reported to have about the same size: but the first appeared significantly lighter than the second (Munsell 6 vs 4.5). Assuming that the black area functions as background in Figure 15f (and hence expands by amodally continuing behind the oval) and as figure in Figure 15a only begs the question of why, then, the functionally large background in Figure 15f is not perceived to be lighter than the oval in Figure 15d. None of these data support the area rule.

The Effect of Size on Lightness in a Bipartite Dome

Let's take the case of a light dome, where one sector is dark, and consider what happens when the dark sector is increased at the expense of the light region. The light region always appears white, and is unaffected by changes in its relative size; but the dark region becomes progressively lighter (Gilchrist, 2002; see solid symbols in Figure 16). In the original anchoring model the first result is expected, but the second is not. The dark region can lighten only by virtue of the area rule: but the area rule only applies when the dark region occupies more than fifty percent of the field. Therefore, the model would predict a flat line for the first three data points in Figure 16.

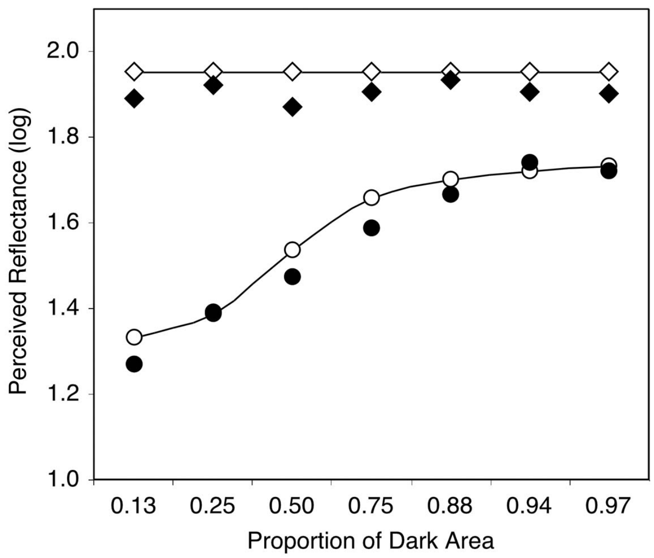

Figure 16. The lightness of the higher-luminance sector in a bipartite dome does not depend on its relative area (solid diamonds), but the lightness of the lower-luminance sector does increase with its relative area (solid circles; Gilchrist, 2002). The open symbols represent the lightness values predicted by the double-anchoring model. See Appendix D.

Note that from a double-anchoring standpoint both regions are functionally surrounds, because neither has fully delimited boundaries. Consider now the dark sector (surround) as it becomes larger and larger. At the surround step, it is always white; at the HL step, its lightness is always equal to its luminance ratio to the light region defined as white. This time, the events of interest take place at the weighting stage. In keeping with the *weight/area* rule we have illustrated earlier, the surround step receives a weight proportional to the area of the surround itself; the HL step, a weight proportional to the area of the highest-luminance region. The light sector is assigned an identical value of white at both steps, so that weighting is inconsequential.

In Figure 16, the open symbols represent the lightness values predicted by the double-anchoring model for two regions whose luminance ratio is 4.39, as in the experiment by Gilchrist (2002). These are the values generated by Equation 2 for *any* two luminances whose ratio is 4.39, when the weight of the surround step is directly proportional, and the HL-step weight inversely proportional, to the relative area of the dark sector.

Testing the Model: Lightness and Darkness Enhancement in Illusory Figures

The original anchoring theory predicts neither the lightness enhancement observed in illusory figures created by black inducers on a white background (as anchors, the illusory figure and the background are expected to look identically white), nor the darkness enhancement observed in their negatives (being judged relative to anchors of the same luminance, the illusory figure and the background are expected to look identically black). In the double-anchoring theory, on the contrary, both effects come as a natural consequence of the surround rule—as illustrated in Figure 17.

In the standard Kanizsa square with black inducers on a white background (top row), for example, the local framework consists in a white illusory region partly surrounded by black inducers. (How the illusory region emerges in the first place is irrelevant here, given that contour formation and lightness enhancement have been shown to be separate phenomena; e.g., Dresp, Lorenceau, & Bonnet, 1990). The illusory region is white at the HL step, and superwhite at the surround step, when it becomes anchored to a white-by-default black surround (the inducers); but the background is white at both steps. In the Kanizsa square with white inducers on a black background (bottom row), on the other hand, the illusory region is black at the HL step and black at the surround step, when it gets anchored to a white surround (the inducers); whereas the background is black at the HL step, but white at the surround step. Thus, the double-anchoring model makes a non-trivial prediction: in the standard display, the illusory *figure* lightens relative to the surround; in the reversed-polarity display, the *surround* lightens relative to the figure. The illusory figure will therefore look lighter than its surround in the white-background display, darker than its surround in the black-background display; and in both cases, the strength of the effect will be a direct function of the contrast of the inducers (as found by Matthews & Welch, 1997).

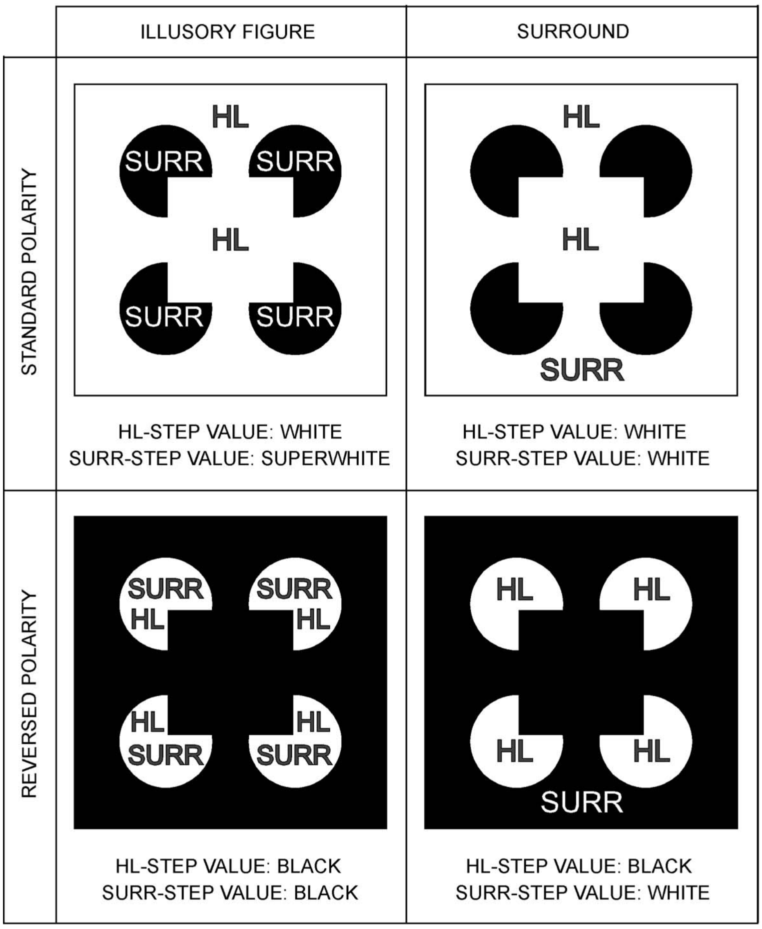

Figure 17. Double-anchoring explanation of lightness (standard-polarity display, top row) and darkness (reversed-polarity display, bottom row) enhancement in the Kanizsa square. Local values at each step are shown for the illusory figure (left column) and the surround (right column). As can be seen, the lightness difference between illusory figure and surround arises, in both displays, at the surround step. HL = highest luminance; SURR = surround.

Testing the Model: "Illumination Discounting"

Influence of Depicted Illumination on Lightness

Altering the simple stimulus of Figure 1 so as to make it more consistent with two evenly illuminated objects, as in the top panel of Figure 18, produces a modest simultaneous contrast effect. But if the scene layout suggests that one of the objects lies in shadow and the other in light, as in the middle panel of Figure 18, the illusion increases considerably. All disks are identical, though, and their local surrounds are also identical.

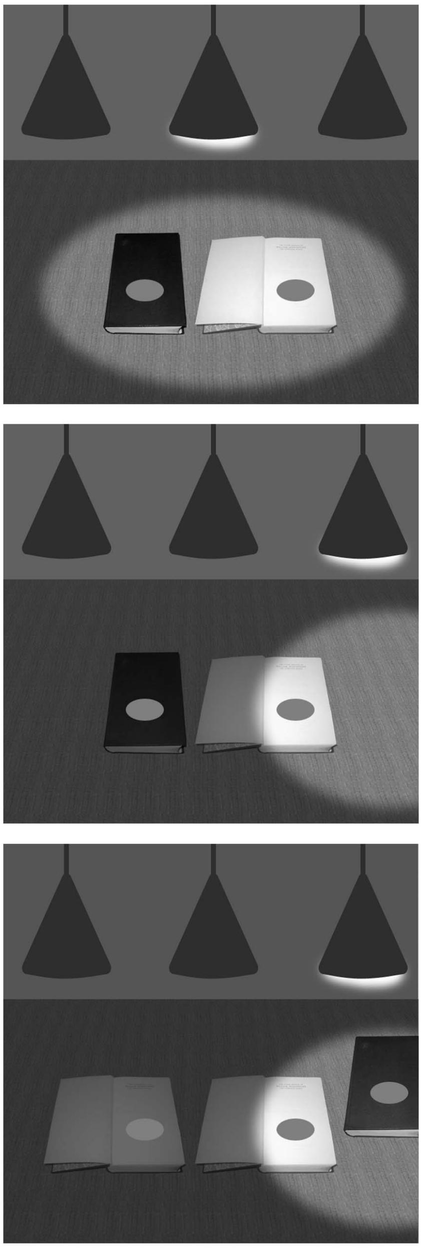

Figure 18. The effect of depicted illumination on lightness. Top panel: The two disks are assigned different lightness values locally but equal lightness values superlocally. Middle panel: The two disks are assigned different lightness values locally and different lightness values superlocally. Bottom panel: The different lightness values assigned superlocally weigh more than the different lightness values assigned locally. All disks are the same shade of gray.

Effects such as these have been used to back up the idea that simultaneous lightness contrast cannot be explained other than "empirically"; that is, on the basis of past *experience* with stimuli generated by different combinations of illumination and reflectance (Purves, Williams, Nundy & Lotto, 2004; Williams, McCoy, & Purves, 1998). Yet, the influence of portrayed illumination on lightness is no real challenge to an anchoring model. In both the top and middle panels of Figure 18 the local frameworks, one black and one white, are equal for each pair of disks, and determine locally identical contrast effects. But in the top panel the two disks belong to the *same* super-local framework (the spotlight), via the grouping principle of common region, and in this shared framework their lightnesses are equal. In the middle panel the two disks belong to *different* super-local frameworks (spotlight and shadow), and in these frameworks their lightnesses are unequal; the disk on black, being the highest luminance, is actually white.

These different super-local assignments carry a lot of weight in the final average, because the spotlight and shadow frameworks group together in an overlay framework (strengthened by top-down information about the scene layout). We thus expect super-local assignments to weigh more than local assignments. This is shown in the bottom panel of Figure 18. Paradoxically, the disk on the open book in the spotlight appears more similar to the disk on the black book in the spotlight (right) than to the disk on the open book in the shadow (left), although local ratios would dictate the opposite. The leftmost disk has a grey surround, but looks distinctly lighter than the rightmost disk, which has a black surround. Again, this might be labeled reverse contrast, but can also simply be interpreted as ordinary contrast in a framework less local, but more powerful.

Lightness of Objects Under Two Illuminations Levels

Another problem that has so far been considered to transcend the scope and explanatory capacity of an anchoring model is that of judging the lightness of a single object that lies half in the shadow and half in the light (Zdravkovic & Gilchrist, 2000). One example is the open book in the middle panel of Figure 18 (disregard the grey disks now). In the double-anchoring model, two illumination levels cast across an object create two adjacent frameworks sustained by luminance polarity and similarity (a portion of the object plus other regions in the shadow, another portion of the object plus other regions in the light), and one overlay framework based on good continuation and conveyed by X-junctions (formed where the contour of the object meets the border between the two regions of illumination). Observers can make separate lightness matches for the shadowed region, the non-shadowed region, and the whole object. For the first two matches, observers are required to pay attention only to one region at a time (the one in the shadow, or the one in the light). This deactivates the overlay framework, which is by definition founded on the extended context, thereby conflating the illuminant and object components.

When asked to estimate the lightness of the object (its “true” color), however, observers must take the whole object and its context into account, establishing the overlay framework of everyday vision. In the model, the lightness of the whole object is simply a weighted average of the lightnesses of its two parts, the one in the shadow and the one in the light—that is of the lightnesses computed in the two frameworks that compose the overlay framework. Such weighting depends on the familiar parameters of size, articulation, and absolute luminance. If one of the two frameworks is larger, and the other is brighter (say, half of the object lies in the same illumination as the rest of the visual field, and the other half is illuminated by a spotlight), we expect a compromise between the two lightnesses. But if one of the two frameworks is larger, more articulated, and more intensely illuminated than the other (say, half of the object lies in the same illumination as the rest of a highly articulated scene, and the other half is under a shadow) the final object lightness will in practice be determined entirely by the lightness computed in the well-lit framework. Both expectations coincide with the results obtained experimentally (Zdravkovic & Gilchrist, 2000).

Snake Illusion and the Shredded and Reversed Snakes

Other examples of overlay frameworks can be seen in Figure 19, which depicts Adelson's (2000) snake illusion together with two of Bressan's (2001) variations (the shredded snake and the reversed snake). In the snake illusion (left display) the top two diamonds look much lighter than the bottom two. These pairs of diamonds belong to an overlay framework of the strong form created by non-reversing X-junctions (i.e., X-shaped junctions where both pairs of edges, left-right and top-bottom, preserve contrast sign), which brings about adjoining stripes of shadow and light. The top diamonds are full increments in their framework; as such, they are white at the HL step and superwhite at the surround step. The bottom diamonds are full decrements in their framework; as such, they are grey at both steps, because they are assessed relative to the white half-ellipses and to their own light grey background respectively.

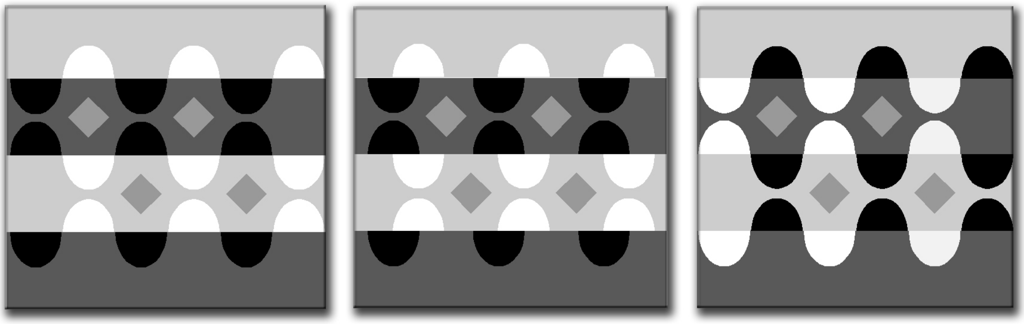

Figure 19. Left panel: The snake illusion (Adelson, 2000). Middle and right panels: The shredded snake and the reversed snake. The contrast effect (whereby the top diamonds appear lighter than the bottom diamonds) is large in the snake, a little weaker in the shredded snake, and virtually gone in the reversed snake. All diamonds are the same shade of gray. The left panel is reprinted with permission from Figure 20, p. 349, of the chapter "Lightness Perception and Lightness Illusions" by E. H. Adelson, as published in M. Gazzaniga (Ed.), *The New Cognitive Neurosciences*, copyright 2000 by MIT Press. The middle and right panels are reprinted with permission from Figure 6, p. 1038, of "Explaining Lightness Illusions," by P. Bressan, as published in *Perception, 30*, copyright 2001 by Pion Limited, London.

In the middle display of Figure 19, I have cut up the snake into horizontal shreds, and shifted every second stripe. This has replaced X-junctions with T-junctions, and destroyed the overlay framework. Here, we expect a lightness illusion similar to that observed in the intact snake (local assignments are the same as in the intact snake) but not as powerful, because, by virtue of their being weaker than the components of overlay frameworks, standard frameworks weigh less in the final average. (The final average includes peripheral assignments, identical for all diamonds.) A slight but statistically significant decrease of the effect is indeed found experimentally (Bressan, 2001).

One might object that overlay frameworks are simply old Helmholtzian "illumination discounting" in new clothing. Objects that appear to sit in the shadow, or under a dark filter, must be, or have repeatedly turned out to be in the past (Purves et al., 2004), lighter than equiluminant objects sitting in the light, and are therefore seen as such. The right display of Figure 19 is at odds with this argument, and shows that the explanatory power of illumination discounting is, at best, inadequate (see also Bressan, 2003). Here, by swapping the external and internal half-ellipses of the snake, I have reversed their polarities. X-junctions are non-reversing, exactly as in the snake, and two adjacent stripes of shadow and light come into view, exactly as in the snake. Illumination discounting should lead to a lightening of the *bottom* diamonds relative to the top diamonds, since, this time, it is they that are covered by a dark medium. Yet, this display generates no contrast illusion at all (Bressan, 2001). From an anchoring perspective the explanation is simple: all diamonds are increments relative to half of their frameworks, but decrements relative to the other half. As such, they are darkened and lightened at the same time.

Testing the Model: Inter- and Intra-Individual Differences

Equivalent Surrounds and Conflicting Frameworks

Imagine a 2x2 checkerboard, and in the middle a target square whose luminance is between the luminances of the checks, as in the bottom rows of Figure 20. Now imagine an identical target square, sitting on a homogeneous surround whose luminance is equal to the checkerboard's average luminance. The original anchoring model requires that the first square appear darker than the second, by virtue of its being assessed relative to a higher local highest luminance (the bright checks). Increasing the checkerboard contrast, all the way to black and white, further augments the highest luminance and should make the square accordingly darker and darker.

It has been shown (Schirillo & Shevell, 1996) that this is indeed true when the target is closer in luminance to the bright than to the dim checks (a case I shall call "partial increment"), but not when it is closer in luminance to the dim than to the bright checks (a case I shall call "partial decrement"). In the former case, the darkening effect is obvious and increases in a very clean manner with checkerboard contrast. In the latter case, different observers behave differently and between-session variability tends to be high for some. Such an asymmetry arises naturally in the double-anchoring model.

The framework arrangement for the two types of target (partial increment, partial decrement) is shown in Figure 20. The bright and dim checks represent opposite contrast polarities relative to the target. This originates two potential (alternative) local frameworks, each containing the target and *one* set of checks, either the bright or the dim set. All conditions are equal here except for luminance similarity, on whose basis partly incremental patches tend to group with the brighter checks (top panel, top row), whereas partly decremental patches tend to group with the dimmer checks (bottom panel, top row). These local frameworks are nested into a super-local framework (target plus checkerboard, bottom row in both panels), founded upon adjacency.

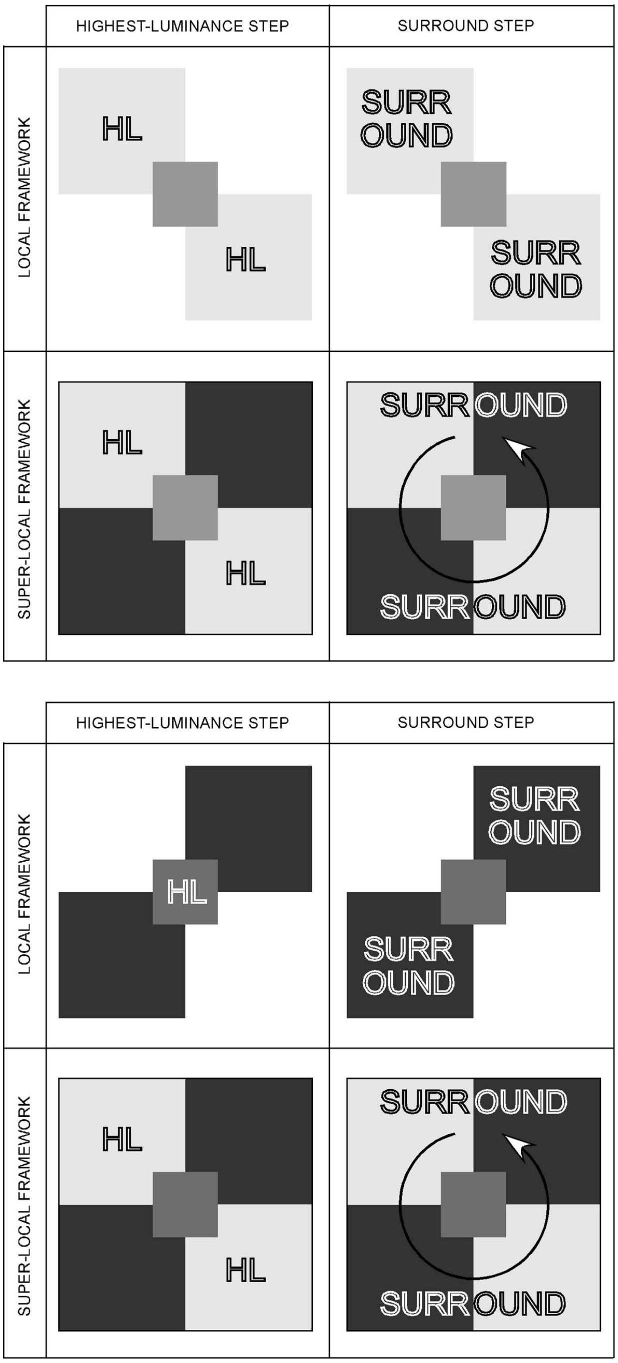

Figure 20. Grouping in partial-increment (top panel) and partial-decrement (bottom panel) checkerboard displays, according to the double-anchoring model. The central square participates in two nested frameworks, local (top row) and superlocal (bottom row). Within each framework, the central square receives two lightness assignments: one at the highest-luminance (HL) step (left column), which is determined by its luminance ratio to the highest luminance, and the other at the surround step (right column), which is determined by its luminance ratio to the surround. In the local framework, the role of surround is played by the more similar checks; in the superlocal framework, the role of surround is played by the average luminance of the checkerboard.

Note that neither framework is stable. Grouping with the checkerboard is strengthened by adjacency, but weakened by conflicting contrast polarities. Grouping with the appropriate set of checks is strengthened by luminance similarity, but weakened by spatial arrangement—that is, discouraged by the specific layout of T-junctions (target and checks do not share the stem of the Ts).

As can be seen in Figure 20, top panel, in the case of partial increments the lightness assignments that the target receives in the two frameworks (relative to an identical target on the luminance-equivalent surround) go in the *same* direction. In the super-local framework, the target darkens, because at the highest-luminance step it is anchored to the brighter check. In the local framework the target darkens even more, because it is anchored to the brighter check both at the highest-luminance and at the surround steps. For partial increments, then, luminance-similarity grouping leads to an increase of the illusion, and darkening is expected to increase neatly with checkerboard contrast.

The same process, however, yields antithetical results for partly decremental patches, which tend to group with the dimmer rather than with the brighter checks (Figure 20, bottom panel). The lightness assignments that the target receives in the two frameworks (relative to an identical target on the luminance-equivalent surround) go in *opposite* directions. In the super-local framework, the target darkens, because at the highest-luminance step it is anchored to the brighter check. Yet in the local framework the target lightens, since it is white at the highest-luminance step (where it is the highest luminance), and superwhite at the surround step (where it is anchored to the dimmer checks defined as white). The final lightness of the target will be a weighted compromise between these two rival tendencies.

Darkening due to super-local grouping is expected to increase with checkerboard contrast, because the luminance of the bright checks increases. However, *lightening* due to local grouping is also expected to increase with checkerboard contrast, the reason being that superwhite induction at the surround step increases—because the luminance of the dimmer checks decreases. The relative strength of the second force, and therefore the final balance, will depend on the weight given to luminance grouping, making for precarious settings. Relatively small weights will yield unstable darkening, whereas under strong luminance grouping the two conflicting forces will cancel each other out, resulting in no darkening whatsoever. Both outcomes have been observed empirically (Schirillo & Shevell, 1996). Interestingly, it can be shown that the tendency of individual subjects to give more or less weight to luminance grouping is consistent across experiments. For a detailed analysis see Bressan (2006).

Modulation of Lightness by Voluntary Attention

As we have seen in our discussion of the Wolff effect, it has been known for a long time that a region can vary in lightness as a function of whether it is interpreted as figure or as ground (Coren, 1969). From a double-anchoring perspective the new effect shown in Figure 21, top display (Tse, 2005), and currently unexplained, falls under the same umbrella. If one looks at the central fixation spot while attending to any of the three grey disks, the attended disk appears to darken. This disk is automatically pulled in front of the others, and becomes an overlapping transparent surface, that is a figure, while the others are forced to play background. As attention shifts from one disk to another, the role of figure swaps accordingly.

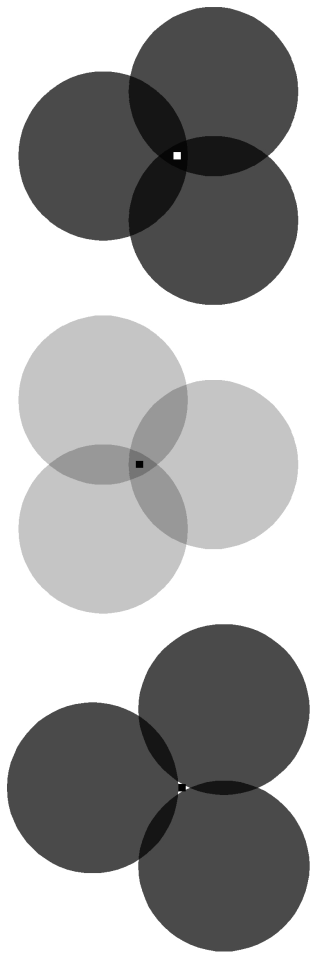

Figure 21. Shifting attention between the disks, while fixating the central point, brings about a change in their lightnesses; the attended disk appears to darken. The top display is similar to the one originally described by Tse (2005). The illusion is weaker in the middle display and gone in the bottom display.

In the model, a decremental disk appears to darken when it is seen as figure relative to when it is seen as background for the same reason why the decremental disks in Wolff's illusion (Figure 11) look darker than the decremental background. In the same vein, if the disks in Figure 21 were increments a symmetrical effect should be observed, whereby the attended disk looks slightly lighter. This is precisely what has been reported (Tse, 2005).

According to the model, the effect is more compelling than Wolff's and Coren's effects because it is founded on overlay, rather than standard, frameworks. This interpretation of the illusion suggests two further predictions. First, if the relative darkening of the front disk is due to the grey of the background disks being diluted by a white assignment at the surround step, the illusion should diminish with very light disks. This can be seen in the middle display. Second, if the effect is based on the attended disk being pulled in front relative to the others in an overlay framework, the illusion should depend crucially on depth/transparency cues such as the area of overlap. This can be seen in the bottom display.

Testing the Model: Summary and Conclusion

In this article, I have shown that an anchoring model of lightness similar to Gilchrist et al.'s (1999) can succeed in predicting a large variety of empirical data, provided that anchoring to highest luminance is complemented by anchoring to surround luminance, and that the definition of frameworks is modified and expanded.

I have shown that the necessity of surround anchoring is independently demonstrated by: (1) the effects of surround luminance (Bressan & Actis-Grosso, 2001), surround articulation (Figure 7; Bressan & Actis-Grosso, in press), and surrounding gradient polarity (the Christmas wall-of-blocks, Figure 2; see also Bressan, 2001) on highest-luminance targets; (2) the effect of remote surround luminance on targets surrounded by identical highest luminances (the butterfly illusion, Figure 8); (3) the effect of relative area on the lightness of the darker sector in a bipartite dome, when this sector occupies less than fifty percent of the field (Gilchrist, 2002); (4) lightness and darkness enhancement in illusory figures. None of the above effects is predicted by highest-luminance anchoring alone, whether or not this is assisted by scale normalization and an area rule.

The necessity of introducing a combination of double anchoring and modified frameworks is demonstrated by: (1) the staircase Gelb Effect: the exact slope of the curve (Cataliotti, 1993), the decrease in compression with room illumination off (Cataliotti, 1993), the difference in curves obtained with and without special illumination in a black laboratory (Bressan, 2002), hypercompression (Gilchrist, personal communication), and insulation (Gilchrist & Cataliotti, 1994); (2) several hitherto unexplained data on lightness in bipartite domes (Li & Gilchrist, 1999); (3) opposite effects between one increment/one decrement and double-decrement versions of the dungeon illusion (Figure 9); (4) inconsistent outcomes of an increase in the highest luminance in equivalent-surround experiments, and inter- and intra-individual differences in some conditions but not others (Schirillo & Shevell, 1996); (5) modulation of lightness by voluntary attention (Tse, 2005). None of the above effects is predicted by the original single-anchoring model.

The present paper, unlike Gilchrist et al.'s (1999), offers quantitative predictions. These are generated by using two simple ratio rules; by using frameworks that are transparently determined by grouping laws whose existence has been demonstrated independently, thereby avoiding a vague or *post hoc* approach; by using weights that are constrained by a conceptually meaningful principle. This principle affirms that a contextual region's size, articulation, and absolute luminance (in nature, the latter covaries with the intensity of illumination) signal how rich in information that region is, and therefore how much faith should be placed in its contribution.

It might be objected that my model fits the empirical data so well just because weights can be adjusted at will. The model uses two weights: the *relative weight* of surround anchoring (i.e., the ratio between the weights of surround and highest-luminance anchoring in the local framework), and the *relative weight* of local anchoring (i.e., the ratio between the weights of the local and peripheral frameworks). This potential objection is invalid for two reasons. First, the weights *are not free parameters*: whereas it is true that these ratios can be appropriately augmented or reduced, whether the numerator is larger or smaller than the denominator is always constrained (for example, an increase in surround size cannot possibly lead to a decrease in surround-anchoring weight). Second, the same weights are used for *all* data points in the data set. This generates just *one* slope, and no creative use of weights can guarantee that such a slope will fit all points.

It might also be objected that my model must necessarily do better than Gilchrist et al.'s (1999) because it has one additional parameter, and the latter will improve the fit because of its mathematical advantage, not of its conceptual validity. This criticism can be rebutted on two accounts—a general one and a circumstantial one. First, the argument would hold if this parameter were a *free* parameter, but surround anchoring is a *fixed* parameter. Even though its contribution is partly adjustable, a wrong fixed parameter would worsen the fit, rather than improve it. Second, and more importantly, it is not at all obvious that my model has an additional parameter. Gilchrist et al.'s scale normalization and area rule are parameters as well. And although it is true that my model has also three weighting factors (relative size, articulation, and absolute luminance), Gilchrist et al.'s model has four (perceived field size, articulation, configuration, and insulation).

Moreover, the choice of surround anchoring as a parameter can be justified on independent grounds, both theoretical and empirical. In the language of the model, a "surround" is any region that groups with the target in a framework. Because the target ultimately groups with *all* regions in the visual field to some degree or other (at least by retinal proximity for lack of stronger forces), surround anchoring simply expresses the notion that the target is influenced by *all* contextual luminances. Such influences must, of course, be scaled by relevance. To recover the "true" grey shade of objects, we need information about the intensity of the light that bathes them. The more it "belongs" to the target, the more a contextual region is likely to share the same illumination, and thereby provide relevant information. This is why the weight of another region's luminance in the assessment of the target's lightness must fall off with decreasing proximity, depth similarity, alignment, and luminance polarity or similarity—i.e., with decreasing likelihood that the two regions are identically illuminated. From an ecological perspective, this is what the concept of grouping or belongingness is all about.

The influence of contextual luminances must be scaled by "relevance of relevance" or information reliability as well: once the extent to which they belong to the target has been estimated (first order of relevance), regions that convey more abundant and dependable information about the illumination level and the target's "true" color should be given more importance (second order of relevance). It is in the interest of accurate, veridical perception that large, richly articulated surrounds are given more weight than small, uniform ones; or that stable, well-lighted frameworks are given more weight than volatile, poorly-lit ones. From an ecological perspective, this is where the weighting rules stem from.

In the model, (a) hierarchical frameworks, and (b) weighting within and between frameworks are simply two aspects of the very same idea, that is two ways to accomplish scaling by relevance. Frameworks express how much each contextual luminance belongs, and—as I have explained—therefore matters, to the target; weighting rules express how much that relationship is going to count relative to all the others.

Rationale of Double Anchoring

In order to be of any practical help in identifying the elements in a scene, the lightness of objects must remain the same whenever the amount of light falling upon them changes, either over time (overall changes in illumination) or over space (local changes in illumination: *illumination-independent*, or Type I constancy). The lightness of objects must also remain the same whenever they are seen against backgrounds of different luminances (*background-independent*, or Type II constancy).

Now, it is easy to see that we do not need double anchoring to solve the problem of lightness constancies. In principle, constancy over time is perfect with either highest-luminance or surround anchoring alone. Both rely on the extraction of a luminance ratio between surfaces in a scene, and on the definition of one of the members of such a ratio (the highest luminance in the first case, the surround in the second) as "white". Because the luminance ratio between the regions in a scene is unaffected by the overall intensity of the illumination, the resulting lightness values remain constant. (In the double-anchoring model, such a constancy is actually expected to be approximate rather than perfect, because of the small but not negligible role of absolute luminance in determining the relative weights of the two anchoring steps.)

Constancy over space, on the other hand, cannot be perfect with either type of anchoring. Imagine two identical objects on a uniform surround, and imagine that the sun shines on half of the scene (one object, and its local surround), and not on the other half. This is a scene that is supposed to give rise to *illumination-independent* constancy. Locally, the two objects are assigned identical values under either type of anchoring. As soon as they get anchored to the peripheral framework, however, constancy is necessarily lost. Under highest-luminance anchoring, this happens because the lower-luminance object becomes anchored to the brightest of the regions in the sun; therefore, it ends up looking darker than the other. Under surround anchoring, on the other hand, the lower-luminance object becomes anchored to the region in the sun, whereas the higher-luminance object becomes anchored to the region in the shade. The first ratio is much smaller than the second, so that, again, the object in the shade ends up looking darker than the other. In short, under either kind of anchoring we expect spatial underconstancy.

Consider now the case where two identical objects stand on two different backgrounds, black and white, under the same illumination. This is a scene that is supposed to give rise to *background-independent* constancy. Under highest-luminance anchoring, the object on black is locally white (whereas the object on white is not), thus the object on black ends up looking lighter than the other. Under surround anchoring, on the other hand, the object on black is locally superwhite (whereas the object on white is not), thus, again, the object on black ends up looking lighter than the other. Under either kind of anchoring we expect simultaneous lightness contrast.

Regarding the problem of lightness constancy, then, two anchorings are not better than one: either will do, admirably when the level of illumination changes over time, imperfectly when luminance variations occur over space, as in Figure 22. Hence one may ask, why should we be equipped with two?

Figure 22. My living room. The lighter of the disks under the chair is the same gray as the darker of the disks in front of the window. The lighter of the disks under the chair appears nearly fluorescent and the lighter of the disks in front of the window does not, although the first has a Munsell value of about 6, whereas the second, with a Munsell value of 9.5, is the highest luminance in the scene. The two disks under the chair are the same gray as the two disks pasted on the mirror above the chair. Impressive illusions such as these are assisted by top-down grouping, which creates here strong overlay frameworks that are interpreted as separate regions of illumination.

Highest-luminance and surround anchoring have in essence the same structure, and generate similar outcomes, with one outstanding exception: the perception of luminosity. I suggest that double anchoring evolved from the necessity of spotting, quickly and reliably, regions of extraordinary luminance. Such are, for example, regions that emit light, such as fire, or reflect it specularly, such as water in sunlight, or predators' eyes at night. To a visual system relying on a surround-as-white principle only, all regions that represent luminance increments relative to their backgrounds would look luminous. To a visual system based purely on a highest-luminance-as-white principle, on the other hand, no surface would ever look luminous. A combination of the two mechanisms permits the discrimination of surface colors from light sources and reflective regions. The latter can be seen as luminous uniquely by virtue of surround anchoring; and can be discriminated from all other incremental surfaces, and therefore pop out, only when anchoring to the surround is assisted by anchoring to the highest luminance.

References

Adelson, E.H. (1993). Perceptual organization and the judgement of brightness. *Science, 262,* 2042-2044.

Adelson, E. H. (2000). Lightness perception and lightness illusions. In M. Gazzaniga (Ed.), *The New Cognitive Neurosciences* (pp 339-351). Cambridge, MA: MIT Press.

Annan, V., Economou, E., & Gilchrist, A. (1998). Locus of error in simultaneous lightness contrast. *Investigative Ophthalmology and Visual Science, 39,* 158.

Arend, L.E., & Goldstein, R. (1987). Simultaneous constancy, lightness, and brightness. *Journal of the Optical Society of America, 4,* 2281-2285.

Arend, L.E., & Spehar, B. (1993). Lightness, brightness, and brightness contrast: 2. Reflectance variation. *Perception & Psychophysics, 54*, 457-468.

Ben-Av, M. B., & Sagi. D. (1995). Perceptual grouping by similarity and proximity: Experimental results can be predicted by intensity autocorrelations. *Vision Research, 35,* 853-866.

Bindman, D., & Chubb, C. (2004). Brightness assimilation in bullseye displays. *Vision Research, 44,* 309-319.

Bonato, F., & Cataliotti, J. (2000). The effects of figure/ground, perceived area, and target saliency on the luminosity threshold. *Perception & Psychophysics, 62,* 341-349.

Bonato, F., & Gilchrist, A.L. (1999). Perceived area and the luminosity threshold. *Perception & Psychophysics, 61,* 786-797.

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

Bressan, P. (2002). [Variations on the Staircase Gelb Effect.] Unpublished raw data.

Bressan, P. (2003). A fair test of the effect of a shadow-incompatible luminance gradient on the simultaneous lightness contrast. Comment. *Perception*, 32, 721-723.

Bressan, P. (2005). The dark shade of the moon. *Clinical and Experimental Ophthalmology*, 33, 574.

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

Bressan, P., & Actis-Grosso, R. (2001). Simultaneous lightness contrast with double increments. *Perception, 30,* 889-897.

Bressan, P., & Actis-Grosso, R. (2006). Simultaneous lightness contrast on plain and articulated surrounds. *Perception*, 35, 445-452.

Bruno, N., Bernardis, P., & Schirillo, J. (1997). Lightness, equivalent backgrounds, and anchoring. *Perception & Psychophysics, 59,* 643-654.

Cataliotti, J. (1993). Distance independent strength of edge integration (Doctoral dissertation, Rutgers University, 1993). *Dissertation Abstracts International, 54,* 03-B.

Cataliotti, J., & Gilchrist, A.L. (1995). Local and global processes in lightness perception. *Perception & Psychophysics, 57,* 125-135.

Coren, S. (1969). Brightness contrast as a function of figure-ground relations. *Journal of Experimental Psychology, 80,* 517-524.

Dal Toè, R. (2002). *La percezione del colore acromatico: Problemi teorici e metodologici.* [The perception of achromatic color: Theoretical and methodological issues.] Unpublished master's thesis, University of Padua, Padova, Italy.

De Weert, C. M. M., & Spillmann, L. (1995). Assimilation: asymmetry between brightness and darkness? *Vision Research, 35*, 1413–1419.

Dresp, B., Lorenceau, J., & Bonnet, C. (1990). Apparent brightness enhancement in the Kanizsa square with and without illusory contour formation. *Perception, 19,* 483-489.

Gelb, A. (1929). Die 'Farbenkonstanz' der Sehdinge. In W.A. Bethe, Von (Ed.) *Handbuch der Normal und Pathologische Psychologie* (Vol. 12, pp. 594-678). Berlin: Springer.

Gilchrist, A.L. (1988). Lightness contrast and failures of constancy: A common explanation. *Perception & Psychophysics, 43,* 415-424.

Gilchrist, A. (2002, October). *Lightness computation in the simplest possible images.* Paper presented at the Trieste Symposium on Perception and Cognition, Trieste, Italy.

Gilchrist, A.L., & Bonato, F. (1995). Anchoring of lightness values in center-surround displays. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 1427-1440.

Gilchrist, A. & Cataliotti, J. (1994). Anchoring of surface lightness with multiple illumination levels. *Investigative Ophthalmology and Visual Science, 35,* S2165.

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.

Gogel, W. C., & Mershon, D. H. (1969). Depth adjacency in simultaneous contrast. *Perception & Psychophysics, 5,* 13-17.

Guadagnucci, L. (2002). *L'effetto della luminanza del target sul contrasto simultaneo di chiarezza.* [The effect of target luminance on simultaneous lightness contrast.] Unpublished master's thesis, University of Padua, Padova, Italy.

Heinemann, E.G. (1955). Simultaneous brightness induction as a function of inducing- and test-field luminances. *Journal of Experimental Psychology, 50,* 89-96.

Howe, P. D. L. (2005). White's effect: Removing the junctions but preserving the strength of the illusion. *Perception, 34,* 557-564.

Kardos, L. (1934). Ding und Schatten [Thing and shadow]. *Zeitschrift für Psychologie, Erg. bd 23*.

Koffka, K. (1935). *Principles of Gestalt psychology.* New York: Harcourt, Brace, & World.

Li, X., & Gilchrist, A.L. (1999). Relative area and relative luminance combine to anchor surface lightness values. *Perception & Psychophysics, 61,* 771-785.

Logvinenko, A.D. (1999). Lightness induction revisited. *Perception, 28,* 803-816.

Logvinenko, A.D., & Ross, D. A. (2005). Adelson's tile and snake illusions: A Helmholtzian type of simultaneous lightness contrast. *Spatial Vision, 18,* 25-27.

Masin, S.C. (2003a). Effects of algebraic and absolute luminance differences on achromatic surface grouping. *Perception, 32,* 615-620.

Masin, S.C. (2003b). On the determinants of surface brightness. *Psychonomic Bulletin & Review, 10,* 220-223.

Matthews, N., & Welch, L. (1997). The effect of inducer polarity and contrast on the perception of illusory figures. *Perception, 26,* 1431-1443.

Newson, L.J. (1958). Some principles governing changes in the apparent lightness of test surfaces isolated from their normal backgrounds. *Quarterly Journal of Experimental Psychology, 10,* 82-95.
Oh, S., & Kim, J. O. (2004).The effects of global grouping laws on surface lightness perception. *Perception & Psychophysics, 66,* 792-799.
Palmer, S. E. (1992). Common region: A new principle of perceptual grouping. *Cognitive Psychology, 24,* 436-447.
Palmer, S. E., Brooks, J. L., & Nelson, R. (2003). When does grouping happen? *Acta Psychologica, 114,* 311-330.
Purves, D., Williams, S. M., Nundy, S., & Lotto, R. B. (2004). Perceiving the intensity of light. *Psychological Review, 111,* 142-158.
Reid, R.C., & Shapley, R. (1988). Brightness induction by local contrast and the spatial dependence of assimilation. *Vision Research, 28,* 115-132.
Schirillo, J.A., & Shevell, S.K. (1996). Brightness contrast from inhomogeneous surrounds. *Vision Research, 36,* 1783-1796
Stewart, E. (1959). The Gelb effect. *Journal of Experimental Psychology, 57,* 235-242.
Todorović, D. (1997). Lightness and junctions. *Perception, 26,* 379-394.
Wallach, H. (1948). Brightness constancy and the nature of achromatic colors. *Journal of Experimental Psychology, 38,* 310-324.
Williams, S. M., McCoy, A. N., & Purves, D. (1998). The influence of depicted illumination on brightness. *Proceedings of the National Academy of Sciences, USA, 95*, 13296–13300.
Wolff, W. (1934). Induzierte Helligkeitsveränderung. *Psychologische Forschung, 20,* 159-194.
Zaidi, Q., Yoshimi, B., Flanigan, N., & Canova, A. (1992). Lateral interactions within color mechanisms in simultaneous induced contrast. *Vision Research, 32,* 1695–1707.
Zavagno, D. (1999). Some new luminance-gradient effects. *Perception, 28,* 835-838.

Appendix A

Predicted Matching Luminances for Simultaneous Lightness Contrast

L_t	L_s	L_h	W_s	W_h	T_l	$\underline{L_h}$	$\underline{T_p}$	W_l	W_p	L_w	L_M
7.95	35.21	35.21	Any	1	0.23	111.9	0.07	0.3	0.7	111.9	13.14
7.95	57.91	57.91	Any	1	0.14	111.9	0.07	0.3	0.7	111.9	10.17
7.95	93.24	93.24	Any	1	0.09	111.9	0.07	0.3	0.7	111.9	8.43
13.80	35.21	35.21	Any	1	0.39	111.9	0.12	0.3	0.7	111.9	22.82
13.80	57.91	57.91	Any	1	0.24	111.9	0.12	0.3	0.7	111.9	17.66
13.80	93.24	93.24	Any	1	0.15	111.9	0.12	0.3	0.7	111.9	14.63
25.07	35.21	35.21	Any	1	0.71	111.9	0.22	0.3	0.7	111.9	41.45
25.07	57.91	57.91	Any	1	0.43	111.9	0.22	0.3	0.7	111.9	32.08
25.07	93.24	93.24	Any	1	0.27	111.9	0.22	0.3	0.7	111.9	26.58

Note. Matching luminance (L_M) of a test square on a 111.9 cd/m^2 background as predicted from the luminance of a comparison decremental square (L_t) on a variable background (L_s), for three luminances of the comparison square and three luminances of the variable background. All labels are defined in the text. The underscored labels $\underline{L_h}$ and $\underline{T_p}$ stand, respectively, for the highest luminance and the territorial lightness in the peripheral framework. Territorial lightnesses are computed as follows, $T_l = (L_t / L_s \text{ x } W_s + L_t / L_h \text{ x } W_h)/(W_s + W_h)$ and $\underline{T_p} = L_t / \underline{L_h}$. The last row of the table contains the predicted final values, $L_M = [(T_l \text{ x } W_l + \underline{T_p} \text{ x } W_p)/(W_l + W_p)] \text{ x } L_w$.

Appendix B

Predicted Matching Reflectances for the Staircase Gelb Effect

R_t	R_s	R_h	W_s	W_h	T_l	$\underline{R_h}$	$\underline{T_p}$	W_l	W_p	R_w	R_M
3	12.0	90	0.50	1	0.11	3	1.00	1	0.35	90	30.37
12	16.5	90	1.00	1	0.43	12	1.00	1	0.35	90	52.02
30	35.5	90	1.00	1	0.59	30	1.00	1	0.35	90	62.61
59	60.0	90	1.00	1	0.82	59	1.00	1	0.35	90	77.96
90	59.0	90	0.50	1	1.18	90	1.00	1	0.35	90	101.7
3	12.0	90	0.10	1	0.05	90	0.03	1	1.75	90	3.64
12	16.5	90	0.20	1	0.23	90	0.13	1	1.75	90	15.24
30	35.5	90	0.20	1	0.42	90	0.33	1	1.75	90	32.79
59	60.0	90	0.20	1	0.71	90	0.66	1	1.75	90	60.79
90	59.0	90	0.10	1	1.05	90	1.00	1	1.75	90	91.56
3	12.0	90	0.50	1	0.11	3	1.00	1	0.00	90	9.50
12	16.5	90	1.00	1	0.43	12	1.00	1	0.00	90	38.73
30	35.5	90	1.00	1	0.59	30	1.00	1	0.00	90	53.03
59	60.0	90	1.00	1	0.82	59	1.00	1	0.00	90	73.75
90	59.0	90	0.50	1	1.18	90	1.00	1	0.00	90	105.8
3	12.0	90	0.10	1	0.05	3	1.00	1	0.00	90	4.77
12	16.5	90	0.20	1	0.23	12	1.00	1	0.00	90	20.91
30	35.5	90	0.20	1	0.42	30	1.00	1	0.00	90	37.68
59	60.0	90	0.20	1	0.71	59	1.00	1	0.00	90	63.92
90	59.0	90	0.10	1	1.05	90	1.00	1	0.00	90	94.30

Note. Top two blocks: Matching reflectance on white (R_M) of each square in a Gelb staircase as predicted from its reflectance and position relative to the other squares, in a large, illuminated, and highly articulated laboratory. First block: spotlight on; second block: spotlight off. In the spotlight-on condition, the contribution of the peripheral framework (W_p), that is of the laboratory scene, is 0.35:1. Removal of the spotlight is reflected by a five-fold decrease in local surround-step weight (W_s), and by a five-fold increase in the weight of the peripheral framework (W_p). Bottom two blocks: Matching reflectance on white (R_M) of each square in a Gelb staircase in a small and entirely black laboratory. The contribution of the peripheral framework (W_p), that is of the laboratory scene, is vanishingly small (here, zero). Third block: spotlight on; fourth block: spotlight off. Removal of the spotlight is reflected by a five-fold decrease in local surround-step weight (W_s).

Appendix C

Predicted Matching Reflectances for the Effects of Area of Figure on Background

L_t	L_s	L_h	W_s	W_h	T_l	$\underline{L_h}$	$\underline{T_p}$	W_l	W_p	R_w	R_M
16	16	25	1	1.5	0.78	539	0.03	1	0.37	90	52.23
16	16	45	1	1.5	0.61	539	0.03	1	0.37	90	41.01
16	16	69	1	1.5	0.54	539	0.03	1	0.37	90	36.14
16	16	89	1	1.5	0.51	539	0.03	1	0.37	90	34.08
16	16	106	1	1.5	0.49	539	0.03	1	0.37	90	32.95
16	16	151	1	1.5	0.46	539	0.03	1	0.37	90	31.18
16	16	25	17	1.0	0.98	539	0.03	1	0.37	90	65.10
16	16	45	17	1.0	0.96	539	0.03	1	0.37	90	64.06
16	16	69	17	1.0	0.96	539	0.03	1	0.37	90	63.61
16	16	89	17	1.0	0.95	539	0.03	1	0.37	90	63.42
16	16	106	17	1.0	0.95	539	0.03	1	0.37	90	63.32
16	16	151	17	1.0	0.95	539	0.03	1	0.37	90	63.15

Note. Matching reflectance on white (R_M) of a 16 cd/m^2 surround of constant size as predicted from the luminance (L_h) and from the area (either large, top block, or 17 times smaller, bottom block), of an incremental square sitting on it. Shrinking the incremental square (highest local luminance) has two consequences: first, it decreases the weight of the HL step (here, from 1.5 to 1); second, it increases the weight of the surround step (here, from 1 to 17). 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 highest peripheral luminance ($\underline{L_h}$). The relative contribution of the peripheral framework (W_p) is here, for both conditions, 0.37:1.

Appendix D

Predicted Matching Reflectances for the Effects of Area in a Bipartite Ganzfeld

R_t	R_s	R_h	W_s	W_h	T_l	R_w	R_M
12	12	52.68	0.125	8.000	0.24	90	21.57
12	12	52.68	0.250	4.000	0.27	90	24.59
12	12	52.68	0.500	2.000	0.38	90	34.40
12	12	52.68	0.750	1.333	0.51	90	45.52
12	12	52.68	0.875	1.142	0.56	90	50.64
12	12	52.68	0.938	1.066	0.59	90	53.01
12	12	52.68	0.969	1.032	0.60	90	54.15

Note. Predicted matching reflectance on white (R_M) of the dark sector (surround) in a bipartite dome, as a function of its relative area. The surround-step weight is directly proportional, and the highest-luminance-step weight is inversely proportional, to the relative area of the dark sector.

Note. Predicted matching reflectance on white (R_M) of the dark sector (surround) in a bipartite dome, as a function of its relative area. The surround-step weight is directly proportional, and the HL-step weight inversely proportional, to the relative area of the dark sector.

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,* 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[6]

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.

---

[6] 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.

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.

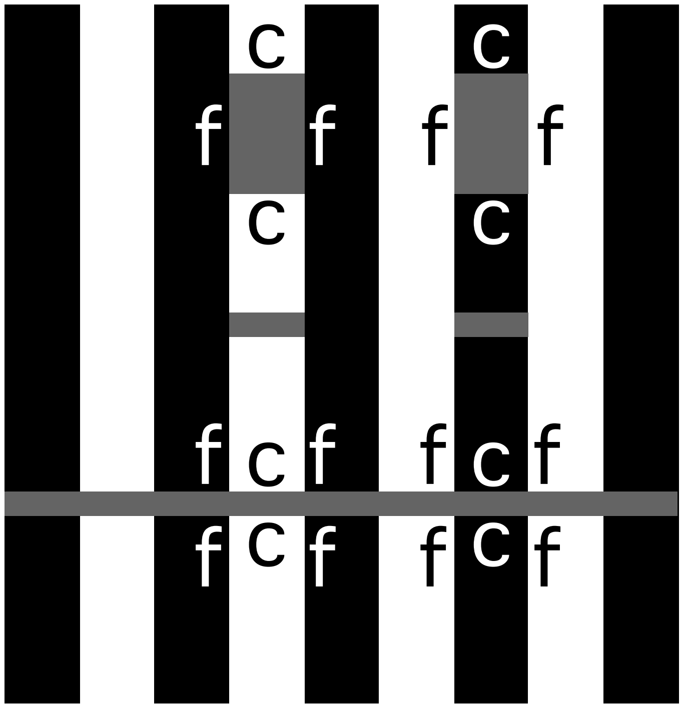

Figure 1. White's effect (top) and grating induction (bottom). In both cases, the gray regions sitting on white look darker than those sitting on black. Each target patch belongs to a collinear-bar framework, *c*, and a flanking-bar framework, *f*. DAT explains either illusion as a combination of anchoring 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.[7]

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.

---

[7] 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.

Bressan, P. (2007). Postscript: The prejudice against frameworks.
Psychological Review, 114, 1115-1115.

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.