Measurements of the effect of surface slant on perceived lightness

Caterina Ripamonti Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA

Marina Bloj Department of Optometry, University of Bradford, Bradford, UK

Robin Hauck Department of Psychology, University of Pennsylvania Philadelphia, PA, USA

Kiran Mitha Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA

Scott Greenwald Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA

Shannon I. Maloney Department of Psychology, University of California, Santa Barbara, CA, USA

David H. Brainard Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA

When a planar object is rotated with respect to a directional light source, the reflected luminance changes. If surface lightness is to be a reliable guide to surface identity, observers must compensate for such changes. To the extent they do, observers are said to be lightness constant. We report data from a lightness matching task that assesses lightness constancy with respect to changes in object slant. On each trial, observers viewed an achromatic standard object and indicated the best match from a palette of 36 grayscale samples. The standard object and the palette were visible simultaneously within an experimental chamber. The chamber illumination was provided from above by a theater stage lamp. The standard objects were uniformly-painted flat cards. Different groups of naïve observers made matches under two sets of instructions. In the *Neutral Instructions*, observers were asked to match the appearance of the standard and palette sample. In the *Paint Instructions*, observers were asked to choose the palette sample that was painted the same as the standard. Several broad conclusions may be drawn from the results. First, data for most observers were neither luminance matches nor lightness constant matches. Second, there were large and reliable individual differences. To characterize these, a constancy index was obtained for each observer by comparing how well the data were accounted for by both luminance matching and lightness constancy. The index could take on values between 0 (luminance matching) and 1 (lightness constancy). Individual observer indices ranged between 0.17 and 0.63 with mean 0.40 and median 0.40. An auxiliary slant-matching experiment rules out variation in perceived slant as the source of the individual variability. Third, the effect of instructions was small compared to the inter-observer variability. Implications of the data for models of lightness perception are discussed.

Keywords: lightness, lightness constancy, scene geometry, surface slant, real objects

Introduction

The perceived lightness of an object provides useful information about the object only if it is stable across the variety of scenes in which the object could appear. A visual system that achieves such invariance is said to exhibit *lightness constancy*. Lightness constancy is difficult to achieve because the light reflected to the observer from an object depends both on the material properties of the object and on the illumination.

Many empirical studies of lightness constancy have investigated how perceived lightness changes with variation in the intensity of a light source (e.g., Henneman, 1935; Wallach, 1948; Arend & Goldstein, 1987), and this is the situation considered in most textbook treatments. It is also possible to manipulate the illumination incident on an object in other ways. Figure 1 shows pictures of the same (planar) object oriented at two different slants with respect to a directional light source. The luminance of the reflected light changes considerably with the change in slant, even

doi:10.1167/4.9.7 Received March 5, 2004; published September 8, 2004 ISSN 1534-7362 © 2004 ARVO

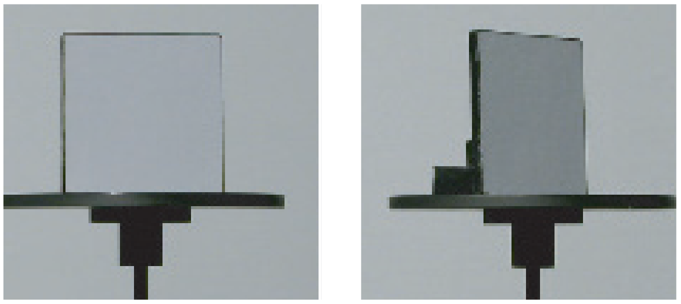

Figure 1. The image shows the same object displayed in two different poses with respect to a light source. The light source is located to the upper left of the object. The luminance reflected from the object to the eye varies with slant.

though the intensity of the physical light source is unchanged.

Mach (Mach, 1886/1959; see Bloj & Hurlbert, 2002) demonstrated that the perceived configuration of a folded card could affect its perceived lightness. Subsequently, a number of studies have demonstrated that perceived geometry interacts with perceived lightness in a manner consistent with lightness constancy (Hochberg & Beck, 1954; Flock & Freedberg, 1970; Gilchrist, 1980; Knill & Kersten, 1991; Pessoa, Mingolla, & Arend, 1996; Williams, McCoy, & Purves, 1998; also Bloj, Kersten, & Hurlbert, 1999; but see Epstein, 1961). These studies are generally qualitative in nature and, although they demonstrate that geometry affects perceived lightness, the range of stimulus configurations where such effects occur and the mechanisms that mediate them are not well understood.

In a recent study, Boyaci, Maloney, and Hersh (2003) report parametric measurements of the dependence of lightness on perceived surface slant in computer-simulated scenes. They argue that the data can be understood within a computational framework, wherein observers estimate and discount the illuminant as a function of the scene geometry. Gilchrist et al. (1999), on the other hand, suggest that effects of geometry are best modeled as arising from (mostly) separate processing within distinct frameworks. Geometry enters the calculation through its effect on how the frameworks are segmented from each other. Adelson (1999) has articulated a similar theoretical perspective.

The present work reports new measurements of how perceived lightness depends on surface slant, and draws some qualitative conclusions. Observers were asked to match the lightness of flat cards displayed at different slants with respect to a single light source. In the companion study (Bloj et al., 2004), we develop a quantitative model of the data.

It is well established that under some viewing conditions different instructions yield different outcomes in matching experiments (e.g., Arend & Reeves, 1986; Bauml, 1999; Bloj & Hurlbert, 2002). Experiments 1 and 2 differed only in the instructions given to observers. In Experiment 1, observers were instructed to match the appearance of the objects. In Experiment 2, observers were instructed to match the paint on the objects.

The data show that for many observers the relation between the luminance of the light reflected to the observer and perceived lightness depends on surface slant, in a manner that tends toward but does not achieve lightness constancy. There are large differences between observers, however, and some observers' data are well described as luminance matches. The variability between observers was large compared to the mean difference induced by our instructional manipulation.

Experiment 3 tested whether inter-observer variability could be attributed to differences in the perception of object slant. It cannot. This experiment also measures how the judgments vary with a change in the position of the physical light source.

Experiment 1

Methods

Apparatus, task, and stimuli

The experimental apparatus consisted of an illuminated booth (see Figure 2). The illumination was provided by a theater stage lamp [SLD Lightning, 6-in. (15-cm) Fresnel #3053, BTL 500-W bulb] placed at the upper left of the booth. Observers viewed stimuli placed in the booth binocularly through a shutter (17 cm wide × 12 cm high) that could be opened and closed under computer control. The observer was positioned inside a separate viewing booth (not shown) that was dimly illuminated by light entering from the back. The observer's head position was stabilized with a chin rest. Information about the light source direction was available to the observer in the form of visible cast shadows (see Figure 2). Observers could not see the light source directly, nor were they explicitly told anything about the illumination in the apparatus. Light reached the stimuli both directly and after inter-reflection within the booth. Information about the relative strength of direct and indirect illumination was potentially available from the shadow contrast.

On each trial of the experiment, observers viewed a *standard object*. The standard object was always a flat card (7 × 7 cm; 3° × 3°) painted matte gray. The card was positioned on a computer-controlled rotatable stage. Eight different standard objects were used. These differed only in the reflectance of the paint. All paints were created by mixing black and white flat acrylic latex house paints (Rich Lux Wal-Shield) in varying proportions. The eight reflectances used were 0.362, 0.265, 0.186, 0.137, 0.114, 0.080, 0.064, and 0.046.[1] The range of standard reflectances was limited to minimize ceiling and floor effects in the matches.

The observer's task was to compare the lightness of the standard object to that of 36 grayscale samples (4 × 4 cm; 1.63° × 1.63°) displayed on a palette that was simultane-

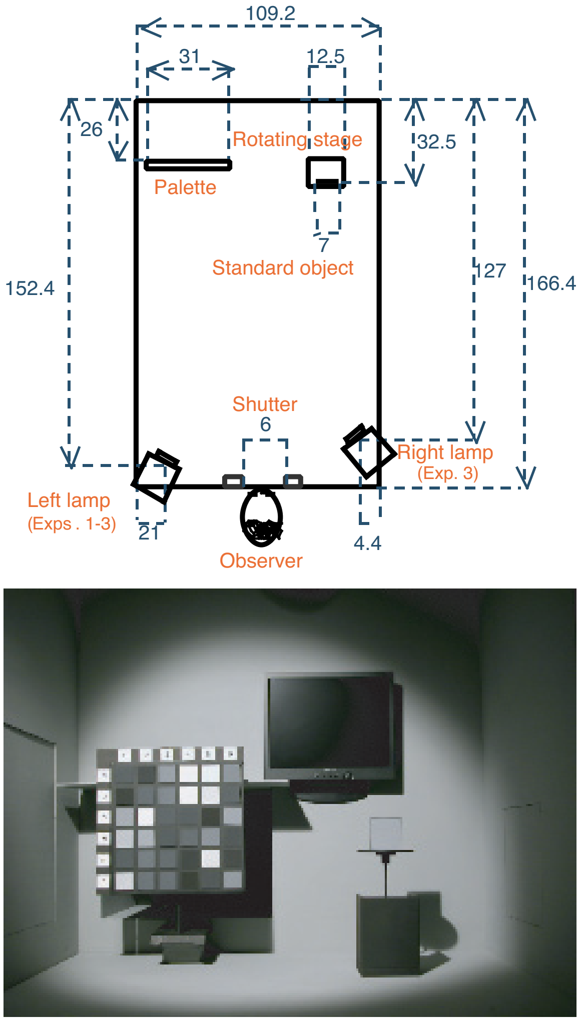

Figure 2. Experimental apparatus. The top panel shows a schematic drawing of the experimental apparatus that consisted of an illuminated booth into which objects could be placed. The azimuth of the left light source was 36° clockwise relative to a reference line normal to the back wall of the booth, while the azimuth of the right light source was 23° counterclockwise with respect to the same reference. The bottom panel illustrates the observer's view of the stimulus. On each trial, both the standard object and the palette were changed through two doors on the sides of the booth. The right light source and the flat-screen monitor seen in the picture were used only in Experiment 3, and were not present during Experiments 1 and 2.

ously visible in the light booth (Figure 2). The samples were arranged on a 6 x 6 grid and the observer chose the palette sample that was closest in lightness to the standard object. The instructions provided to observers are described in more detail below.

The observer responded verbally by reporting the row and column of the selected sample using a letter/number coordinate system (see Figure 2). An experimenter recorded the observer's response on each trial. Observers were allowed to look back and forth between the standard object and the palette for as long as they required. If the observer was unable to find an acceptable match because of gamut limitations on the palette (i.e., there was no surface light or dark enough), he or she was asked to state this rather than choosing the best available sample. This occurred rarely. Across all three experiments reported in this study, the number of trials on which observers were unable to find matches, or on which they chose either the highest or lowest palette reflectance, was less than 1% of the total number of trials (33 trials out of 5460). Data from such trials were excluded from the quantitative analyses described in the "Appendix."

Each palette sample was matte gray, painted using the same technique as the standard objects. The lightest sample had a reflectance of 0.869, while the darkest sample had a reflectance of 0.029. The rest of the samples were selected by the authors to produce an approximately uniform lightness scale. Supplementary material (click on the link to view) tabulates the palette reflectances. All eight standard object reflectances were contained in the palette.

We constructed four separate palettes, each containing the same 36 sample reflectances in a different random arrangement. On each trial of the experiment, a randomly chosen palette was placed in the booth. Each palette could be placed with a different edge at the top, so that there were effectively 16 different random arrangements used in the experiment. We used randomly arranged palettes to discourage observers from employing a strategy where they attempted to produce consistent responses by classifying the standard objects and remembering the palette location they had selected for a given object on a previous trial.

The interior surfaces of the light booth were painted gray using a mixture of the same black and white paints used to paint the samples. The reflectance of these surfaces was 0.375. In addition to the standard object and palette, other objects were visible in the apparatus. These varied from experiment to experiment, but always included the stands that supported the standard object and palettes. For Experiment 3, an LCD flat panel monitor was mounted on the back wall of the booth to enable measurement of the perceived slant of the standard objects. Although it is shown in Figure 2, this monitor was not present for Experiments 1 and 2.

Observers

Seven naïve observers with normal or corrected-to-normal vision participated in the experiment. Observers were screened using two visual tests: Keystone Vision-Screener II to ascertain their visual acuity and stereopsis and Ishihara's test for color deficiency to ascertain their color vision. Observers with corrected Snellen acuity less than 20/30 and/or a stereopsis of less than 75% on the Sheperd-Fry scale were excluded from the study as were those that made any mistakes in the Ishihara test.

Experimental conditions

The standard objects were displayed in five different slants: +60°, +45°, +30°, 0°, and −45°. The 0° slant corresponds to the standard object displayed at an orientation parallel to the back wall of the apparatus. Positive orientations correspond to counter-clockwise rotations when the standard object is viewed from above. An experimental session consisted of 40 trials presented in random order (eight standard object reflectances crossed with five slants). At the start and finish of each trial, the shutter was closed so that the observer could not see into the booth.

Each observer participated in three experimental sessions, for a total of 120 matches per observer.

At the end of the third session, observers were asked to fill in an evaluation form consisting of questions about perceptual strategies used to accomplish the task, comments on the stimuli, etc.

Instructions

Observers were told that during the experiment they would see a series of cards displayed at different slants and that their task on each trial was to identify the palette sample that appeared the same shade of gray as the card. However, observers were allowed to interpret the word "appearance" themselves: They were not explicitly instructed as to whether they should try to match the reflectance of the card and sample or the luminance of the reflected light. For this reason, we refer to these instructions as the *Neutral Instructions*.

The instructions were accompanied by a demonstration in which the experimenter showed a card rotated with respect to a directional light source.

In Experiment 2 below, we used different instructions where observers were explicitly told to judge card and sample reflectance (*Paint Instructions*).

Calibration

The independent stimulus variables in our experiment are the reflectance and slant of the standard objects, while the dependent variable is the palette sample reflectance. Slant was measured using a 1° scale with an indicator needle, mounted permanently on the rotatable stage. The needle was not visible to the observer. Standard and palette reflectances were measured as described above.

In addition to these variables, the actual stimulus seen by the observer depended on the intensity and location of the light source. We were particularly interested in knowing how the luminance reflected from the standard object varied with its slant. At the start of the experiment, this was measured in situ for each standard object reflectance. For all slants except for +60°, the reflected luminance was measured directly using a PhotoResearch PR-650 spectral radiometer placed at the observer's location. For the +60° slant, the visual angle subtended by the standard object was too small to be measured directly by the PR-650. To obtain luminances for this card, we used the PR-650 to calibrate the nominal pixel values of dark-corrected images acquired through a 550-nm interference filter using a high-quality digital CCD camera (Photometrics PXL). The calibration information was then used to infer the luminance of the standard objects at +60° from images taken at this slant. The reflectances, luminances, and chromaticities of the standards are tabulated in the supplementary material. Auxiliary measurements were made on a daily basis to monitor against apparatus drift.

Results

The top panel of Figure 3 shows the matching data for two observers. The mean match reflectance (averaged over the three sessions) is plotted as a function of the slant of the standard object. The individual curves in each plot show data for one standard object reflectance. The vertical shifts between the individual curves indicate an increase of match reflectance with standard object reflectance, as one would expect. It is also clear from the raw data that the observers' reflectance matches depend on the standard object slant.

The different magnitude of the matches for different standards makes it difficult to compare the effect of slant on match reflectance across the standards. In addition we found, consistent with Weber's Law, that the variability of the matches was roughly proportional to the value of the raw match. This effect is illustrated in Figure 4. To visualize the effect of slant across standard reflectances, and to plot

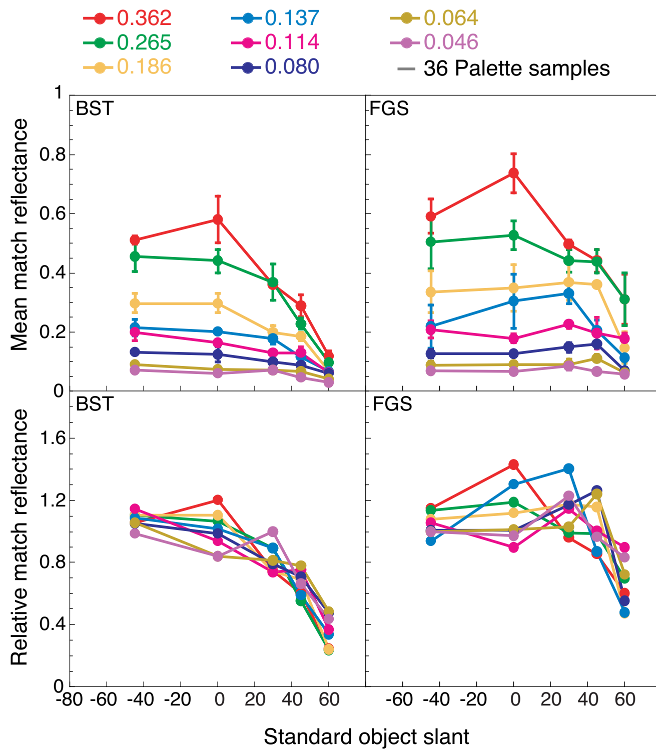

Figure 3. The figure shows mean matching data for two observers. The top panel illustrates the raw data and their respective standard errors for the eight standard objects plotted as a function of slant. The ticks next to the right top panel illustrate the 36 palette samples reflectances; colored ticks illustrate the subset of palette reflectances used for the standard. The bottom panels show the data for each standard reflectance normalized to a mean of 1. Error bars are omitted from the bottom panels for clarity.

the data from different standards on a scale that equalized the precision of the matches across standards, we scaled the mean match reflectance $r_{i,\theta}$ for each standard reflectance i and slant θ to obtain normalized matches $r_{i,\theta}{}^{norm}$:

$$r_{i,\theta}{}^{norm} = r_{i,\theta} / \overline{r}_i , \quad (1)$$

where $\overline{r}_i$ is the mean of the $r_{i,\theta}$ taken over θ. The normalized matches for same two observers are plotted in the bottom panels of Figure 3.

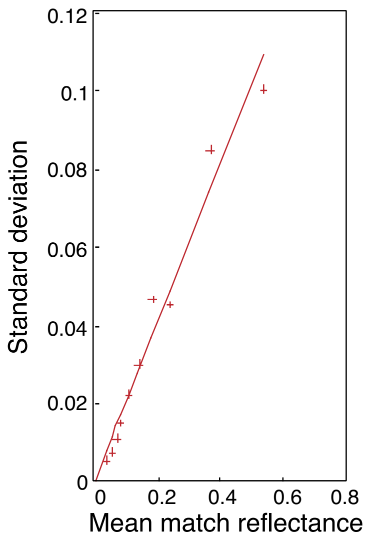

Figure 4. Average standard deviation of matches plotted against corresponding mean match. We divided the range of mean reflectance matches into 10 bins and obtained the standard deviations of the matches (over the three sessions) for each mean match within the bin. We then averaged the standard deviations for all matches within each bin. This average standard deviation is plotted in the figure. The data were aggregated over all standards, slants, and observers from Experiment 1.

The normalized matches show that the effect of slant is similar for each standard reflectance and emphasizes the dependence of the matches on the slant of the standard object. Also clear in the plots is that the detailed dependence of the matches on slant differs between the two observers.

Because our primary interest is in understanding how the appearance of the standard objects varies with their slant, we aggregated the data over standard object reflectance. This aggregation increases our experimental power to address questions about the overall effect of slant. We refer to the aggregated data for each observer at slant θ as the normalized relative match reflectance $\overline{r}_\theta{}^{norm}$. The procedure used to obtain the $\overline{r}_\theta{}^{norm}$ is described in the "Appendix." A feature of the aggregation procedure is that for each observer the $\overline{r}_\theta{}^{norm}$ has a mean over θ of 1. As described in the "Appendix," the procedure also removes from the data of each observer the effect of sigmoidal nonlinearities in the relation between standard and match reflectance.

The differences between observers can be seen even more clearly in Figure 5, where data aggregated over reflectance are shown for all seven observers. The difference in the effect of slant for observers BST and FGS remains quite clear, and the other observers show results that, qualitatively, fall between those for BST and FGS.

The differences between the observers are not due to measurement variability. The error bars in Figure 5 indicate 90% confidence intervals obtained using a resampling procedure.[2] Where error bars are not visible, they are smaller than the plotted points. The 90% confidence intervals are small compared to the differences between observers. In addition, an ANOVA indicates that the differences between observers are statistically significant. The main effect of slant was significant at $p < .001$, whereas the slant by ob-

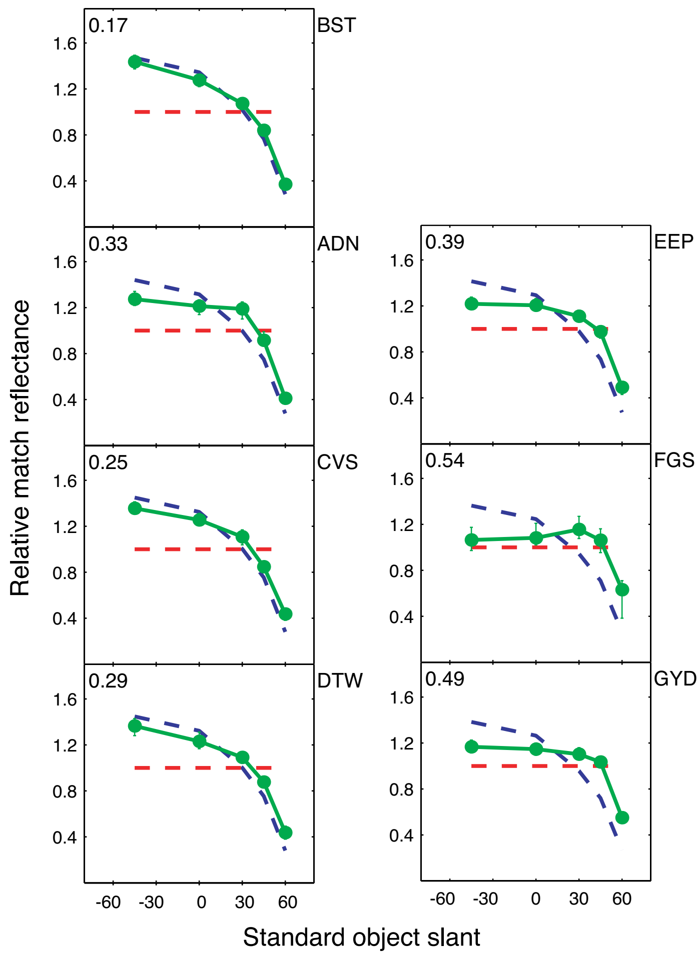

Figure 5. The figure illustrates data for all seven observers who were given Neutral Instructions. Observers' normalized relative match reflectance (connected green symbols) has been plotted versus the standard object slant. The red horizontal broken lines represent the predictions of lightness constancy. The blue broken lines represent the predictions for luminance matches. Error bars show 90% confidence intervals. Observers are listed from top to bottom, roughly in order of increasing constancy, as evaluated by the model-based constancy index developed in the companion work (Bloj et al., 2004). The values at the top left of each panel are the error-based constancy indices for each observer.

server interaction was significant at $p < .01$. Because the data are normalized for each observer, there is no main effect of observer.

The supplementary material tabulates the raw data for Experiment 1, as well as for Experiments 2 and 3 below.

Relation to lightness constancy

The data shown in Figure 5 do not indicate lightness constancy. If the observers' matches were lightness constant, each standard would always be matched with the same palette reflectance, leading to normalized relative match versus slant curves that consisted of horizontal lines (i.e., $\bar{r}_\theta^{norm} = 1$, red broken line in each panel). None of the observers' data are well described by a horizontal line.

Another useful comparison for observer performance is the prediction obtained by assuming that observers make their matches in proportion to the luminance reflected from the standard (luminance matching). In this case, the data should be fit by a scaled version standard card luminance as a function of slant. The predictions for luminance matching are shown in Figure 5 as blue, broken lines. For each observer, the predicted curve has been scaled to provide the best fit to that observer's data.

Observer BST comes close to exhibiting luminance matching. The other observers all show deviations from luminance matching in the general direction of lightness constancy. The data from observers FGS and GYD come close to being constant over a range of standard object slants (–45° to +45°) but deviate substantially as the slant increases to +60°. Other observers exhibit behavior that varied between these extrema.

To provide a sense of how observers vary between luminance matching and lightness constancy, we derived a simple error-based constancy index. We found for each of the three sessions, the normalized relative matches for that session. Let ε_{lum}^2 represent the sum of squared errors between the luminance matching prediction (blue broken lines in Figure 5) and the individual session data for a single observer. Similarly, let ε_{const}^2 be the sum of squared errors for the constancy prediction (red broken lines in Figure 5). Then

$$CI_e = \varepsilon_{lum} / \left(\varepsilon_{lum} + \varepsilon_{const} \right) \quad (2)$$

is a constancy index that takes on a value of 0 when the data are perfectly characterized as luminance matches and a value of 1 when the data reveal perfect constancy. For Experiment 1, the index ranges between 0.17 and 0.54, with a mean of 0.35 and median of 0.33. Table 1 provides the index values for all experiments reported in this paper. Error-based constancy indices for each observer are provided in Figure 5. Interpretation of the index should be tempered against the observation that the index definition is somewhat arbitrary and that no single number can capture the richness of the data. In the companion study, we present model-based summary measures of the individual variation.

The mean data (across observers) from Experiment 1 are plotted in Figure 7 below. This plot also reveals performance that is intermediate between luminance matching and lightness constancy, but it should again be emphasized that individual variation around the mean is large.

Experiment 2

Methods

Apparatus and stimuli

We used the same apparatus and stimuli as in Experiment 1.

Observers

Seven naïve observers who did not participate in Experiment 1 took part in Experiment 2. Observers were screened using the same visual tests and exclusion criteria as in Experiment 1.

Experimental conditions

Experimental conditions were the same as for Experiment 1.

Instructions

Observers were told that during the experiment they would see a series of cards that had been painted using different paints. They were also told that the cards would be displayed at different slants. These instructions were accompanied by a demonstration that showed a card rotating under a directional light source. The rotation of the card served to show that, in the demonstration at least, the apparent lightness changed across slants. Observers were then told that their task was to pick the sample from the palette

Error Based CI	Minimum	Maximum	Mean	Median
Exp 1	0.17	0.54	0.35	0.33
Exp 2	0.24	0.63	0.48	0.52
Exp 3 Left Neutral	0.21	0.45	0.35	0.38
Exp 3 Right Neutral	0.31	0.54	0.43	0.42
Exp 3 Left Paint	0.19	0.51	0.35	0.33
Exp 3 Right Paint	0.36	0.61	0.46	0.45
Overall	0.17	0.63	0.40	0.40

Table 1. Error-based constancy index. The table provides the minimum, maximum, mean, and median values of the error-based constancy index (Equation 2) for all experiments reported in this work.

that had been painted with the same paint as the standard card, despite the fact that the standard might appear different when seen under different lighting. This is essentially the "paper match" task of Arend and Reeves (1986), and we refer to these instructions as *Paint Instructions*.

Calibration

Experiments 1 and 2 were run concurrently, so that the measurements of standard luminance and apparatus stability were common to the two experiments.

Results

Figure 6 illustrates normalized relative match reflectances for all seven observers plotted versus the standard object slant. Normalized relative match reflectances and 90% confidence intervals were obtained using the same procedures described for Experiment 1, and constancy and luminance matching predictions are again shown (cf., Figure 5).

As with Experiment 1, the data from Experiment 2 indicate substantial variability among observers. In an ANOVA, the slant by observer interaction was significant at $p < .001$. Observer LEF's data (Experiment 2) is very close to the luminance matching prediction, just as with observer BST (Experiment 1). Several observers (e.g., JPL, KIR, and NMR) in Experiment 2 performed similarly to observer FGS from Experiment 1. Thus the range of performance observed in Experiment 2 is similar to that observed in Experiment 1. Comparison of Figures 5 and 6, however, does give the impression that on the whole more luminance matching behavior was seen in Experiment 1 and more lightness constant behavior (at least over a range of slants) was seen in Experiment 2. Error-based constancy index values are provided in Table 1.

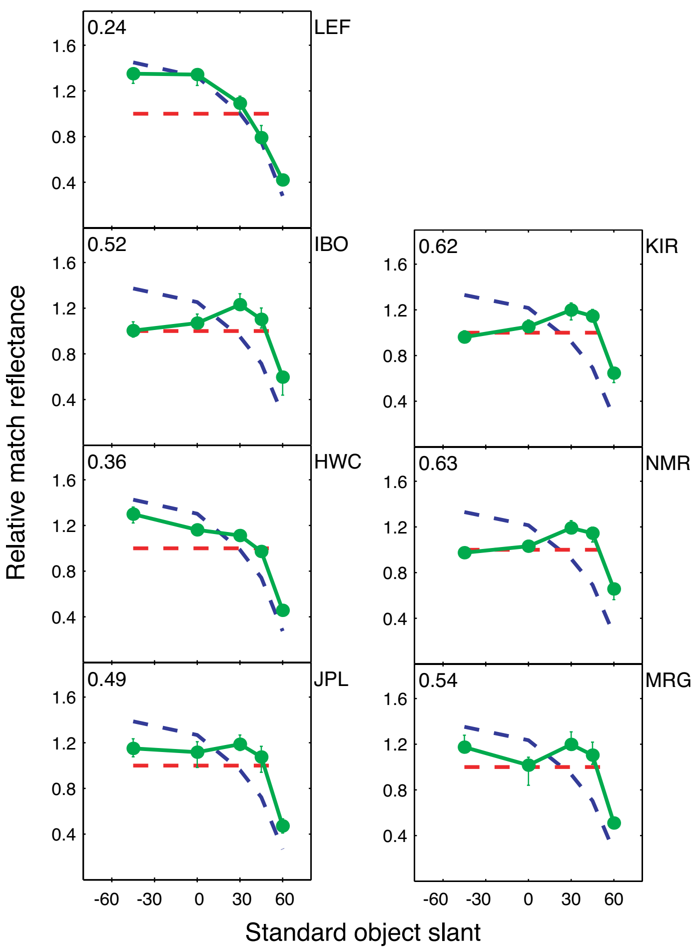

Figure 6. The figure illustrates data for all seven observers who were given Paint Instructions. Same format as Figure 5.

Figure 7 shows the mean data from both Experiments 1 and 2, obtained by averaging the individual normalized matches from each experiment. The difference in means is small compared to the range of performance observed in either experiment, although an ANOVA reveals that it is statistically significant (slant by instructions interaction, $p < .05$).

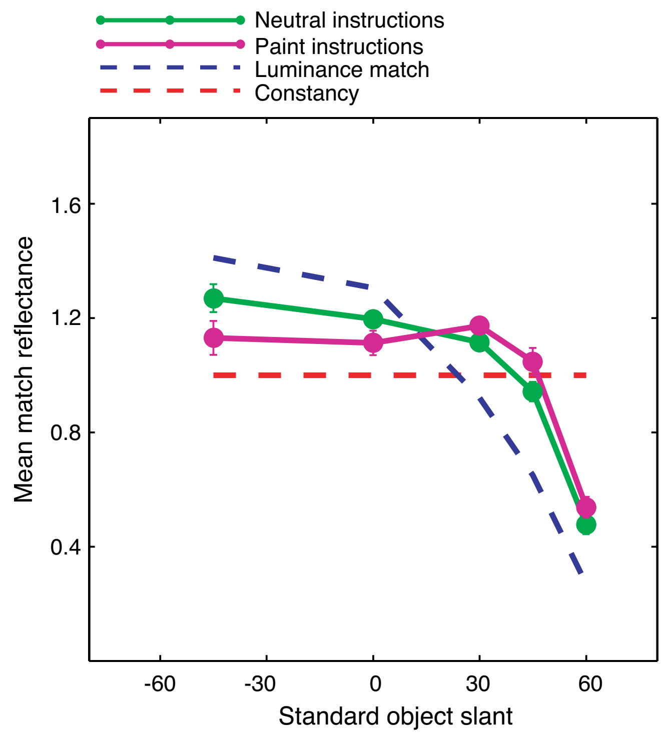

Figure 7. Mean data from Experiments 1 and 2, obtained by averaging the data from individual observers shown in Figure 5 Experiment 1) and Figure 6 (Experiment 2). Error bars show +/- one *SEM*. The predictions for lightness constancy and luminance matching (scaled to the mean of the two plotted curves) are also shown.

Intermediate discussion

Experiment 2 confirms the essential conclusions of Experiment 1: Most observers exhibit performance that lies somewhere between luminance matching and lightness constancy, and there are large individual differences between observers. The effect of instructions is small compared to the range of performance within each instructional group.

One possible cause for the individual differences is that they are indeed due to differences in strategy employed by individual observers, but that our instructional manipulation was not powerful enough to have a decisive influence on the strategy employed by any given observer. Another possibility is that the individual differences have a different origin. In Experiments 1 and 2, we did not independently assess how observers perceived the slant of the standard object. Systematic differences in perceived slant could produce differences in how perceived lightness depends on physical slant, even if the basic mechanisms that integrate slant and lightness are performing identically for all subjects. In Experiment 3, we added a slant-matching task to our protocol to test this idea.

Experiment 3

As noted above, one aim of Experiment 3 was to investigate whether the variability we found in Experiments 1 and 2 could be attributed to variation in the perception of object slant. In addition, we wanted to explore the effect of changing the light source position and to replicate the results of our instructional manipulation.

The experiment consisted of two parts. During the first part, 14 naïve observers were asked to perform lightness matches, using essentially the same procedure employed for Experiments 1 and 2. During the second part, the same observers were asked to match the slant of the standard object. Seven observers were given the Neutral Instructions and seven were given the Paint Instructions.

Methods

Apparatus

We used the same apparatus as in Experiments 1 and 2. The illumination was provided by either of two theater-stage lamps placed above the booth (Figure 2, top panel). One of the two light sources was located at the upper left of the booth, whereas the other was located at the upper right. We set the intensity of the two lamps such that when the standard object was displayed at 0° slant, its luminance would be the same under the two light sources.

Tasks

Observers were required to perform two tasks. The first task was the lightness-matching task used in Experiments 1 and 2. The only change in this task was that rather than reporting their responses verbally, observers used a joystick to control the position of a small red indicator dot projected onto the palette. When the dot was on their preferred palette samples, observers pressed a button to record their choice. This modification eliminated the need for the experimenters to manually enter the observers' verbal responses.

After the lightness-matching task was completed (three sessions), observers performed a slant-matching task. We adopted the task described by Van Ee and Erkelens (1995). Observers adjusted the orientation of a line displayed on an LCD flat panel monitor located on the back wall of the viewing booth (see Figure 2). The image displayed on the monitor provided a schematic representation of the experimental chamber viewed from above (see Figure 8). The image contained a rectangular frame (22 cm wide × 16.5 cm tall; 7.56° × 5.68°) representing the booth, with an aperture on the front which indicated schematically the observer's viewing position. Two lines (11.5 cm; 3.96°) were drawn inside the rectangular frame: One represented a reference line parallel to the back wall of the booth, while the other was the match line controlled by the observers. The lines were drawn using an anti-aliasing algorithm. The orientation of the match line was controlled through two game-pad push-buttons. Observers could also control the angular step size of the rotation (1° or 10°). Observers were asked to adjust the match line so that it provided a top view of the slant of the standard object. The starting orientation of the match line was selected randomly on each trial. On a small number of trials, the observer accepted a match before he or she meant to and indicated this to the experimenter before the start of the next trial. Data from such trials were excluded from the analysis. The raw slant-matching data are available in the supplementary material.

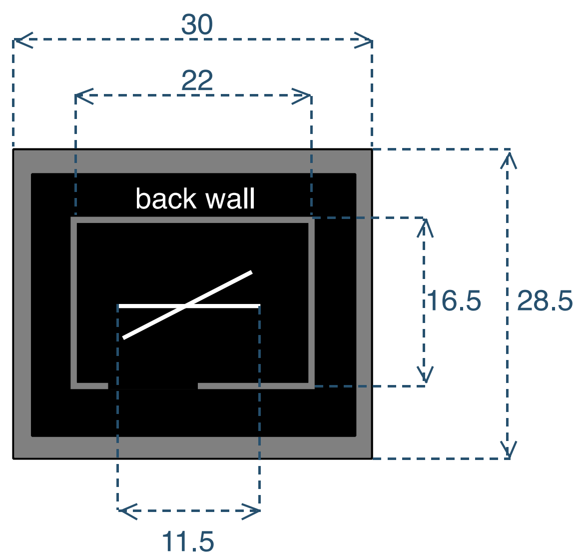

Figure 8. Schematic representation of the slant-matching display. The rectangular frame represented the booth viewed from above. The line parallel to the back wall was held fixed throughout the whole experiment, while the oblique line was adjusted by the observer to match the slant of the standard object. The spatial resolution of the monitor was 832 x 624 pixels. The distance between the monitor and the observer was 161.4 cm.

Observers

Two groups of seven naïve observers participated in the experiment. None of them had participated in either Experiment 1 or Experiment 2. Observers were screened using the same visual tests and exclusion criteria used in Experiments 1 and 2.

Experimental conditions

Lightness matching. During the lightness-matching part, the standard objects were displayed in nine different slants: −60°, −45°, −30°, −15°, 0°, +15°, +30°, +45°, and +60°. Five standard objects were used. As in Experiments 1 and 2, these were flat matte cards. The card reflectances used in Experiment 3 were 0.265, 0.186, 0.137, 0.114, and 0.080.

On every trial, the light source position (left or right) was randomly selected. While the shutter occluded the observer's view of the chamber, both light sources were turned on. Then one of the lights was turned off, leaving the selected light on. Thus the lights cycled between every trial, even if the selected light was the same for two or more consecutive trials.

An experimental session consisted of 90 trials (five standard object reflectances crossed with nine slants and two light sources.) These trials were presented in random order. Each observer participated in three lightness-matching sessions, for a total of 270 matches per observer.

Within a session, the 90 trials were run in two blocks of 45 trials each. The same two blocks were used for all observers and sessions, but the order of the trials within each block was randomly chosen for each observer/session. Observers were allowed to take a 10-min break between blocks.

Slant matching. In the slant-matching part, only two of the five standard object reflectances (reflectances 0.186 and 0.114) were used. The standard objects were displayed in the same nine slants used for the lightness matches. The same two light source positions were also used.

An experimental session consisted of 36 trials (two standard object reflectances crossed with nine slants and two light sources). These trials were presented in random order.

Each observer repeated the slant matches three times in a single session, for a total of 108 matches per observer.

Debriefing. At the end of the slant estimation session, observers were asked to fill in an evaluation form consisting of questions about potential strategies used to accomplish the two tasks, comments on the stimuli, etc.

Instructions

Lightness matching. The 14 observers were divided in two groups. One group was given Neutral Instructions (as in Experiment 1), and the other was given Paint Instructions (as in Experiment 2).

Slant matching. All 14 observers were told that during this session they would see a standard object displayed at various slants and that their task was to adjust the slant of a line on a computer screen such that it matched the standard object's slant. Observers were told that the image displayed on the screen represented a schematic view of the standard object seen from above and that the horizontal reference line on the screen indicated a slant parallel to the back of the apparatus.

Calibration

The same calibration procedures employed in Experiments 1 and 2 were used. Reflectances, luminances, and chromaticities are tabulated in the supplementary material.

Results

Lightness matches as a function of physical slant. Figure 9 illustrates the data for the Neutral Instructions and the left light source position. The plots in the left column show the normalized relative match reflectance as a function of slant. As in Experiments 1 and 2, observers' data exhibit performance that varies from luminance matching (e.g., FGP) toward lightness constancy (e.g., ALR). The data are somewhat noisier than in Experiment 1, a point to which we return below. Error-based constancy index values are provided in Table 1.

Figure 10 shows the data for the same seven observers aggregated over trials where the light source was on the right: generally speaking, the match reflectance decreases with slant when the light is on the left but increases with slant when the light is on the right. This qualitative dependence is what one would expect for observers who perform luminance matches, as can be seen from how the luminance matching prediction changes when the light source is moved. For observers whose performance is approximately characterized as luminance matching for both light source positions (e.g., FGP), it is possible that the effect of changing light source position can be explained entirely in terms of the change in reflected luminance.

Figures 11 and 12 show the data for those observers who were given Paint Instructions. Results for this group of observers are quite similar to those obtained for the Neutral Instructions.

Lightness matches as a function of matched slant. One of the purposes of Experiment 3 was to test whether differences between observers could be attributed to differences in perception of slant. During the second part of Experiment 3, observers were asked to match the slant of the standard objects.

Figure 13 shows the mean slant matches of those observers who were given Neutral Instructions, separated by light source position. Figure 14 shows the same data for observers who were given Paint Instructions. Generally, observers matched slant quite accurately, with only a small tendency for underestimation of absolute slant for slants near 0°. The right columns of Figures 9, 10, 11, and 12 replot individual observers' match reflectances as a function of match, rather than physical slant. This representation does little to reconcile the large individual variation.

Effect of instructions. Figure 15 compares the mean data from each instructional condition. As with the comparison between Experiments 1 and 2, any overall effect of instructions is small compared to the range of performance shown by individual observers. ANOVAs indicated no significant effect of instructions (slant by instructions interaction $p = .87$ for illuminant from left; $p = .75$ for illuminant from right).

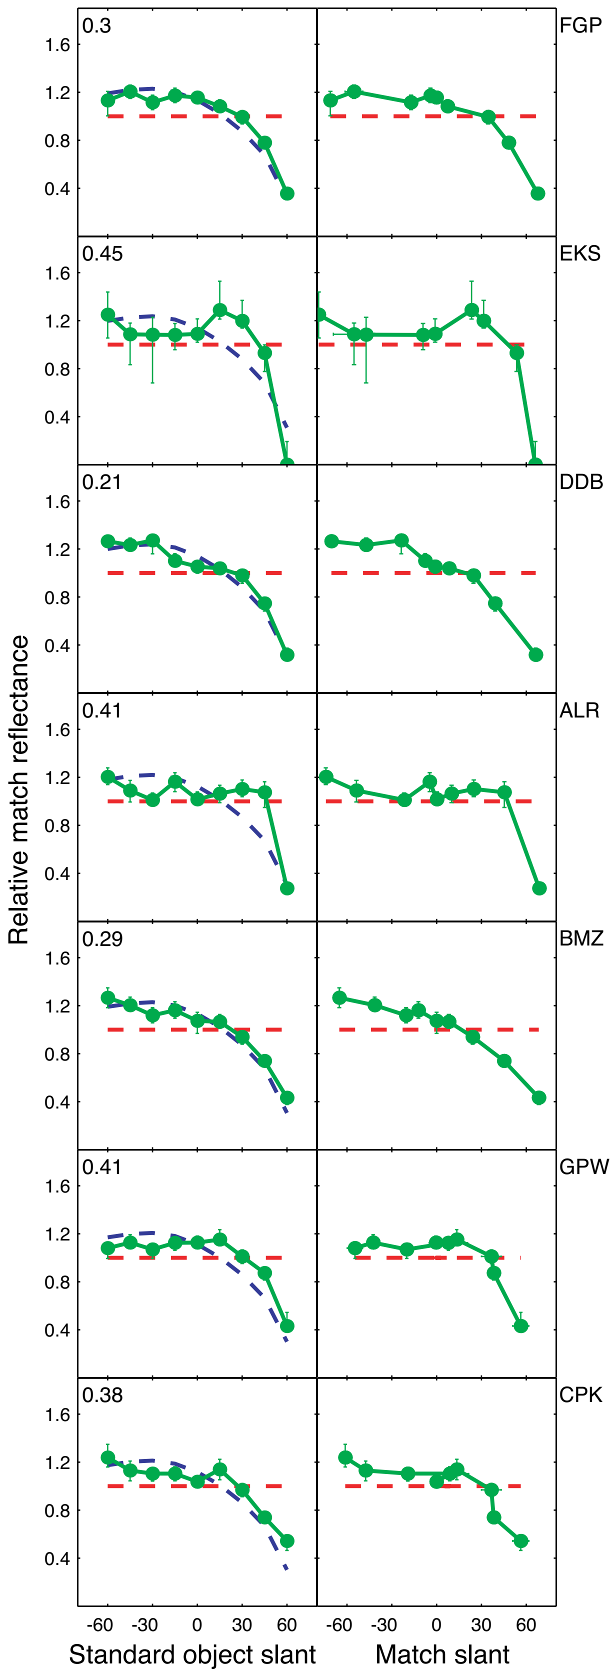

Figure 9. The figure plots normalized relative match reflectances for the trials on which the standard objects were illuminated from the left. These observers were given Neutral Instructions. Left column: matches plotted against physical slant. The red horizontal broken lines represent the predictions of lightness constancy. The blue broken lines represent the predictions for luminance matches. The error-based constancy index for each observer is provided at the top left of each panel. Right column: matches for each observer plotted against the match slant obtained for that observer. Prediction for luminance matches not shown as we did not measure card luminances at observers' matched slants. Error bars show 90% confidence intervals, obtained as described for Experiment 1.

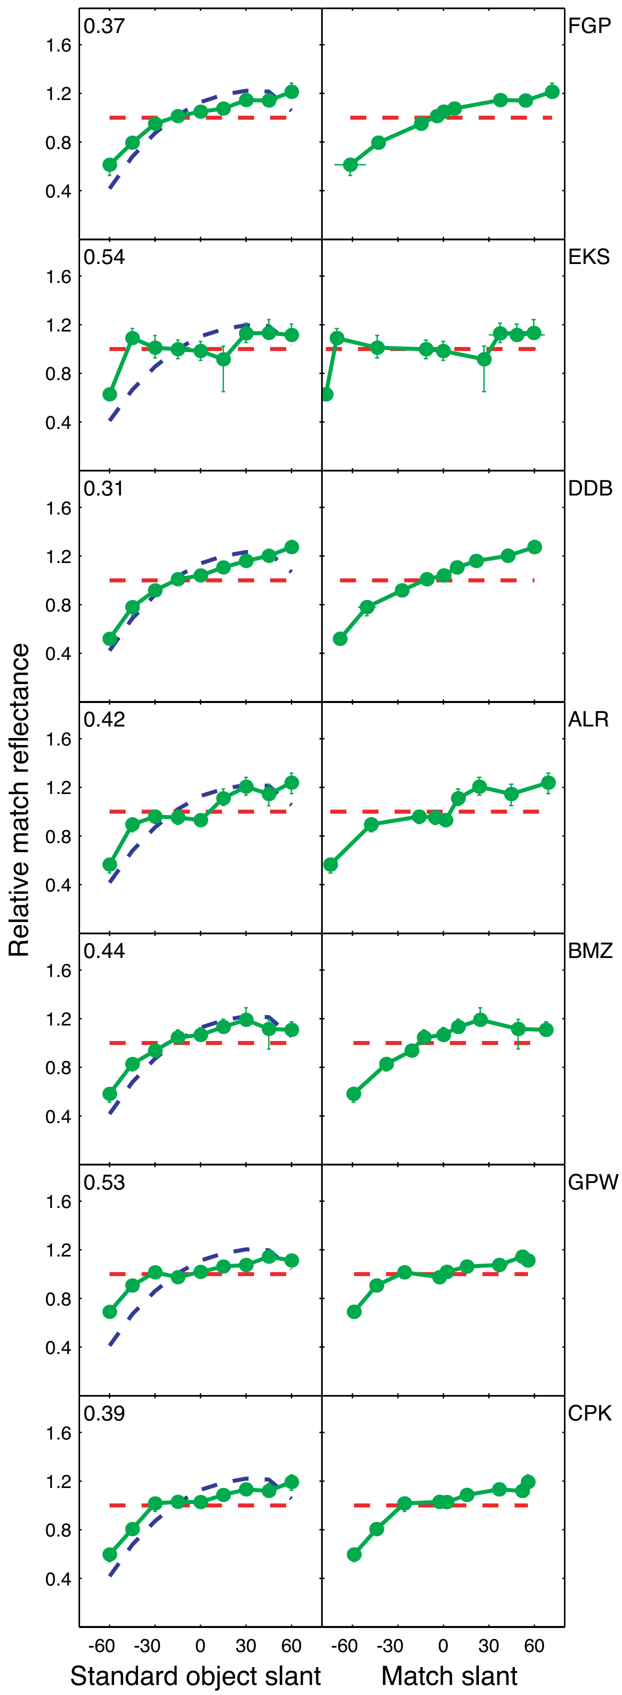

Figure 10. Normalized relative matches for trials on which the standard objects were illuminated from the right; same observers, instructions, and format as Figure 9.

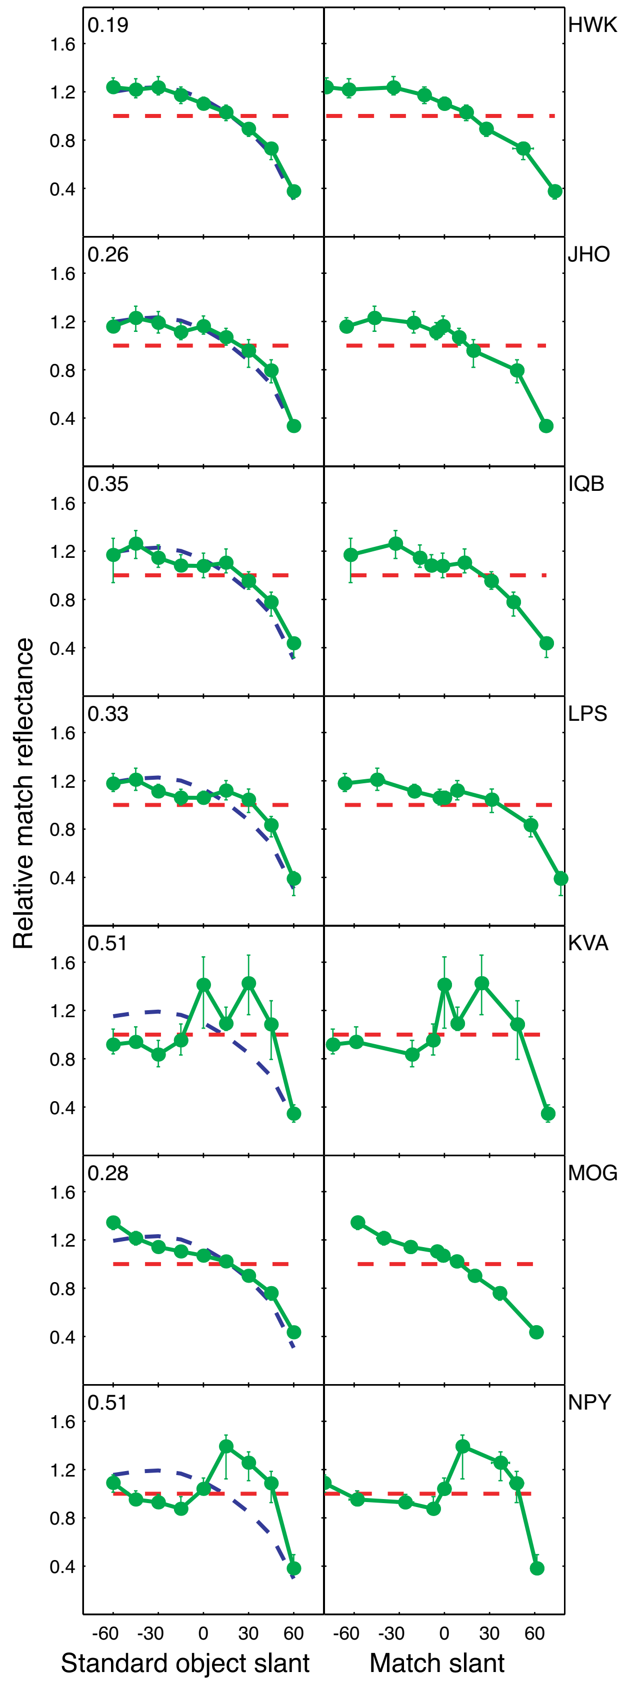

Figure 11. Same as Figure 9 for observers who were given Paint Instructions.

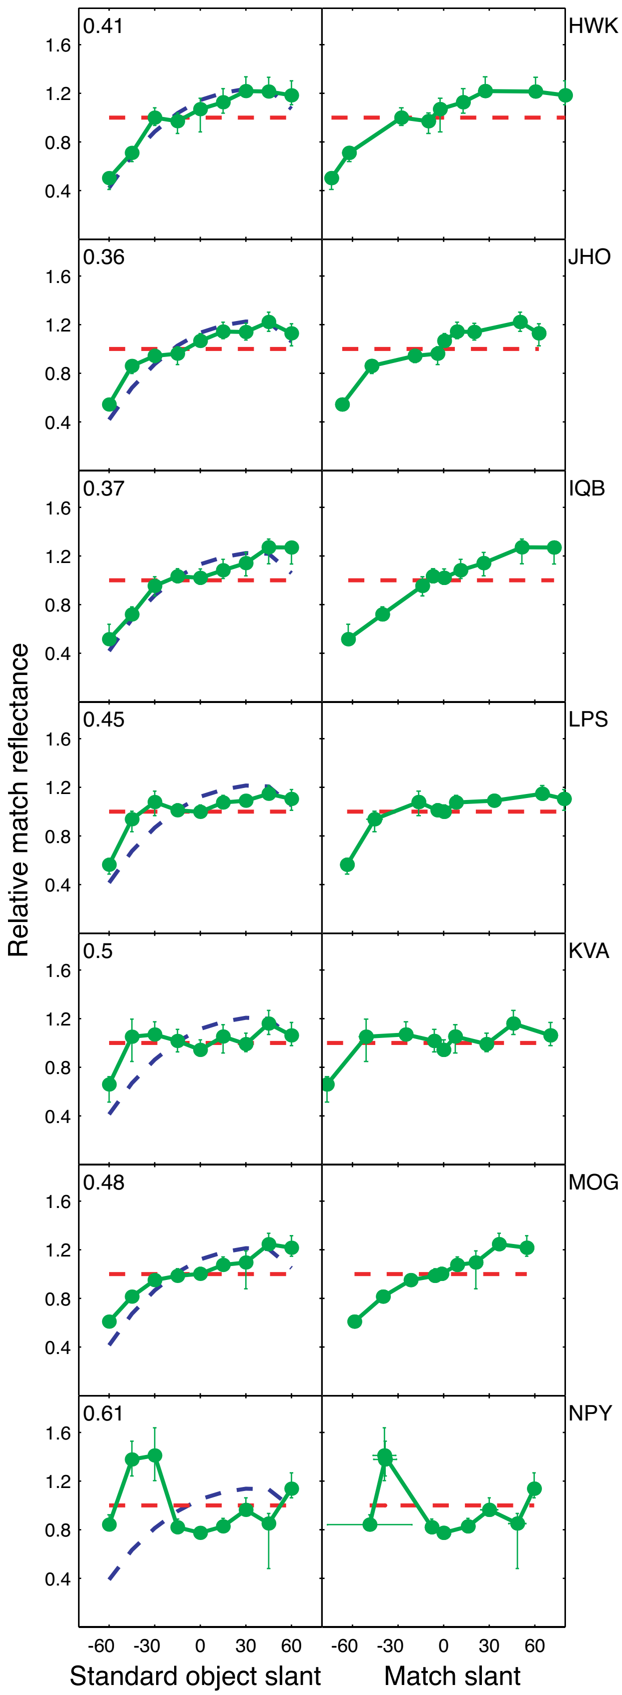

Figure 12. Same as Figure 10 for observers who were given Paint Instructions.

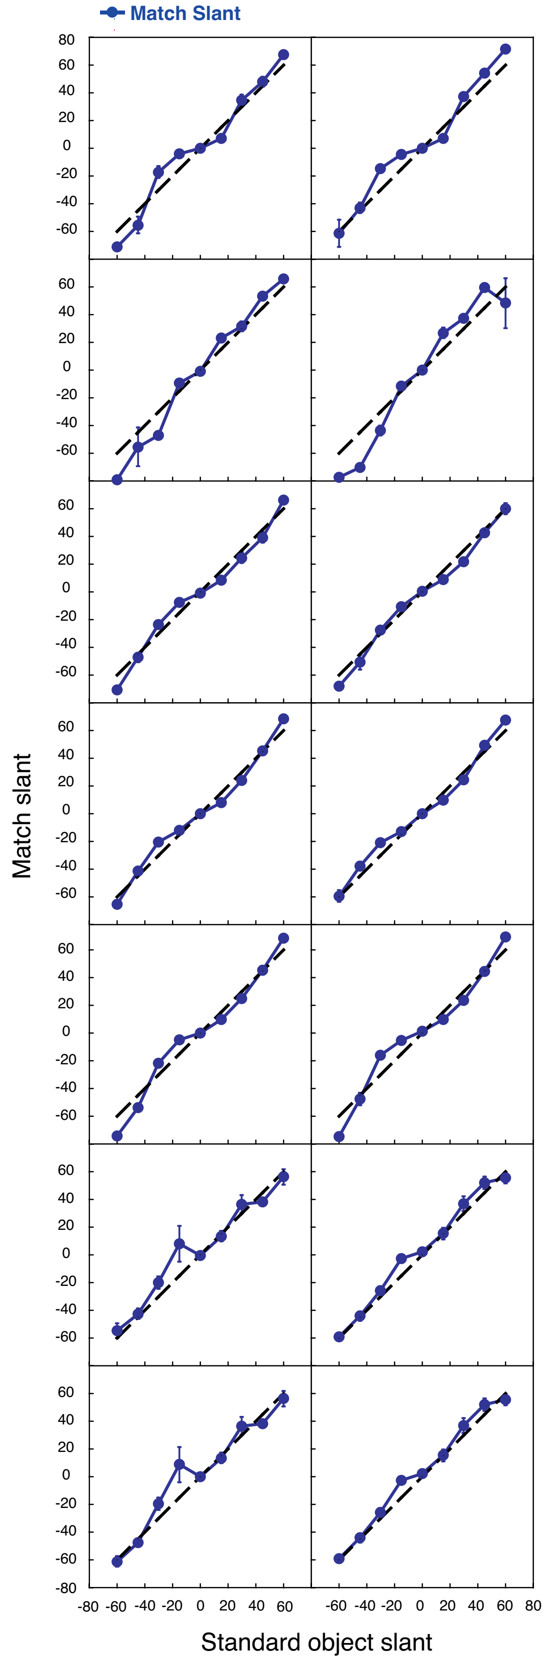

Figure 13. The figure illustrates mean slant matches for observers given Neutral Instructions. Left column: slant matches for illumination from left. Right column: slant matches for illumination from right. Data were averaged over replications and standard object reflectances. Error bars show +/–1 *SEM*.

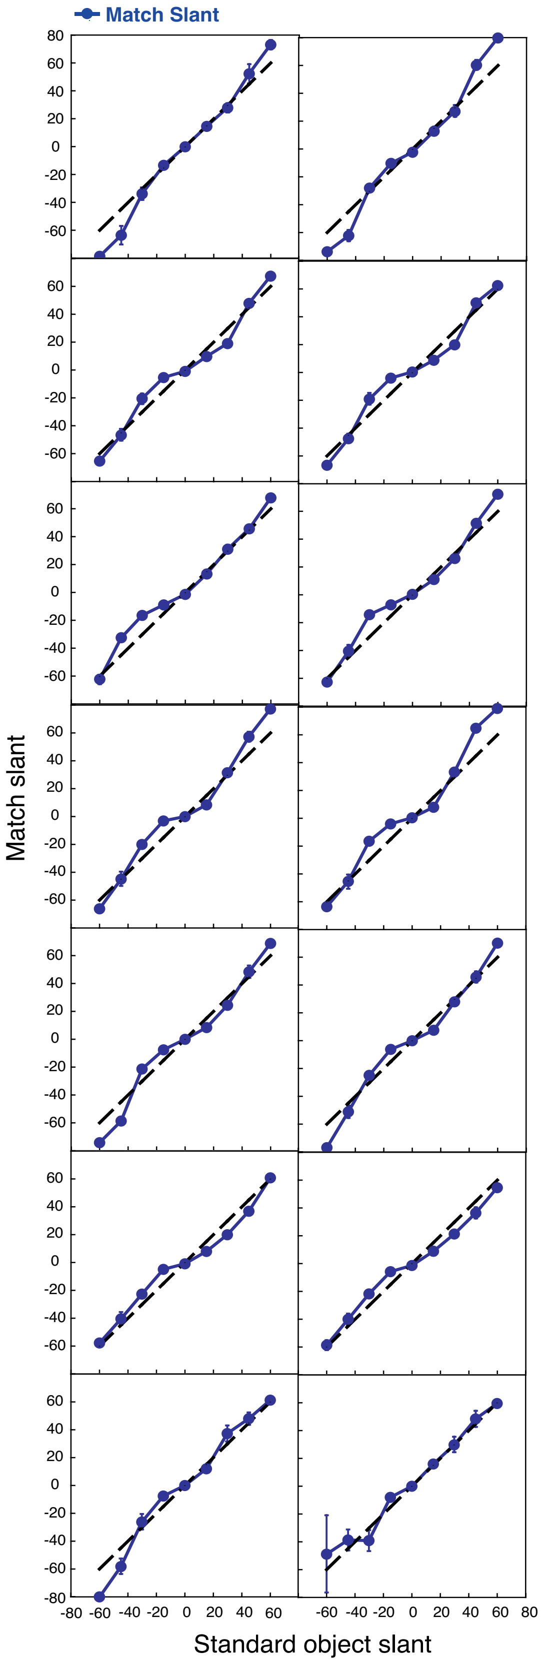

Figure 14. The figure illustrates mean slant matches for observers given Paint Instructions. Same format as Figure 13.

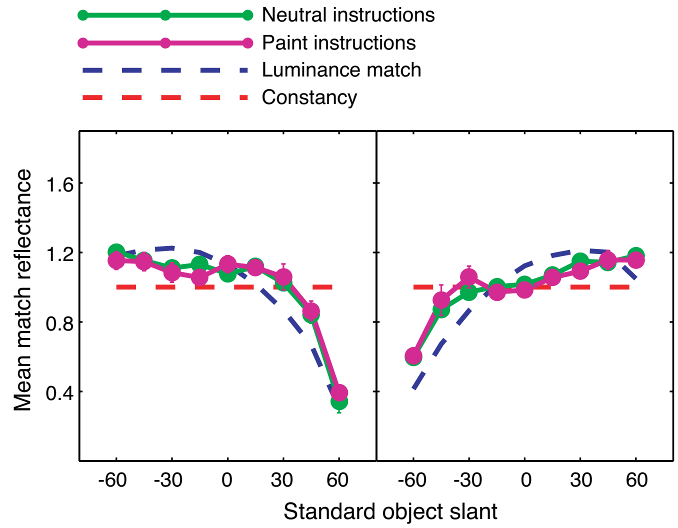

Figure 15. Mean data from Experiment 3. The left panel shows data for trials when the illumination was from the left, while the right panel shows the data when the illumination was from the right. Each panel is in the same format as Figure 7.

Replication of Experiments 1 and 2. The left illuminant position trials of Experiment 3 provide a replication of Experiments 1 and 2, albeit with the right illuminant position trials intermixed. In addition, Experiment 3 employed more slants and fewer standard object reflectances. Figure 16 compares the mean reflectance matches across experiments. The differences are small. For the comparison between Experiments 1 and 3, an ANOVA indicated that

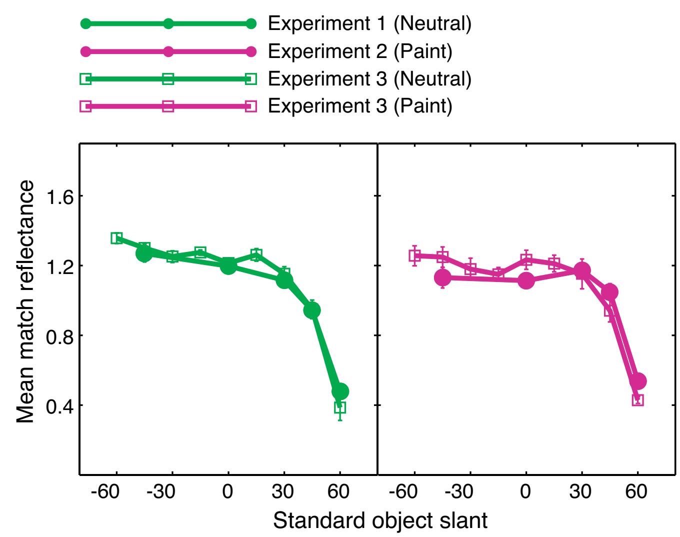

Figure 16. Comparison across experiments. The left panel compares the mean reflectance match obtained in Experiment 3 for the left illuminant position and Neutral Instructions with the mean reflectance match from Experiment 1. The right panel compares the mean reflectance match obtained in Experiment 3 for the left illuminant position and Paint Instructions with the mean reflectance match from Experiment 2. In each panel, the data from Experiment 3 are scaled to have a mean of 1 for the subset of slants for which data were obtained in Experiments 1 and 2. Error bars show +/-1 *SEM*.

the interaction between experiment and slant was not significant ($p = 0.60$), while for the comparison between Experiments 2 and 3, the same interaction was significant at the 0.05 level ($p = 0.052$). The small difference may arise because in Experiment 3 trials with the right illuminant position were intermixed with those with the left illuminant position. Individual observer differences were significant in all four (light source by instructions) conditions of Experiment 3 ($p < .01$).

Variability in Experiment 3. The data from Experiment 3 appear somewhat noisier than the data from Experiments 1 and 2. One possible reason for this is that fewer points were collected for each slant. In Experiments 1 and 2, data for each slant are aggregated over eight standard reflectances, while in Experiment 3, only five standard reflectances were used. A second important difference is that two light source positions were interleaved in Experiment 3. If observers failed to track (either implicitly or explicitly) the change in illuminant position on some trials, this would show up as noise in the data. We have not examined the difference in variability in any formal way, but it may be advantageous to block trials by illuminant position in future experiments.

Discussion

We report experiments that measure how perceived lightness depends on slant. Consistent with most of the earlier literature (Mach, 1886/1959; Hochberg & Beck, 1954; Flock & Freedberg, 1970; Gilchrist, 1980; Williams et al., 1998; Bloj et al., 1999; Boyaci et al., 2003), our data indicate that the visual system takes the scene geometry into account as it computes object lightness. Most of the earlier studies report only data averaged over observers, and certainly the average observers in our experiments show performance that deviates systematically from luminance matching toward lightness constancy (see Figures 7 and 15).

Individual variability

A striking feature of our data is the large and reliable individual differences. Although the average observer clearly takes geometry into account, some of our individual observers exhibit performance that is close to that predicted by simple luminance matching, while others show a considerable degree of lightness constancy. It is not possible to evaluate the degree of individual observer differences from most previous reports. An exception is a recent study by Boyaci et al. (2003). Our design was very similar to theirs, with the important exception that they used synthetically generated stimuli displayed on computer-controlled monitors. Figure 17 replots their individual observer data in the same general format as Figure 5. As with our observers, most of their observers show performance that is intermediate between luminance matching and lightness constancy, although by eye their data are closer to luminance matches than ours. There also appear to be systematic differences between their individual observers, although again these seem smaller than the ones we measure. It is possible that differences between their experiments and ours are due to effects of the use of synthetic versus real images.

The slant-matching data in Experiment 3 rule out the possibility that the individual differences arise because of individuals' disparate perceptions of slant. There was considerable consistency in the slant matches across observers, and replotting the lightness data as a function of individually matched slant did not reduce the variability.

It remains possible that different observers bring different strategies to the lightness matching task, and that the differences between observers result from different strategic choices. For example, some observers may be trying to match standard object luminance while others attempt to match standard object reflectance. Our instructional manipulation failed to separate observers into two groups. This may be because our Neutral Instructions were weaker than the appearance instructions that have been used previously (e.g., Arend & Reeves, 1986; Bauml, 1999; Bloj & Hurlbert, 2002). It remains of interest to pursue this issue

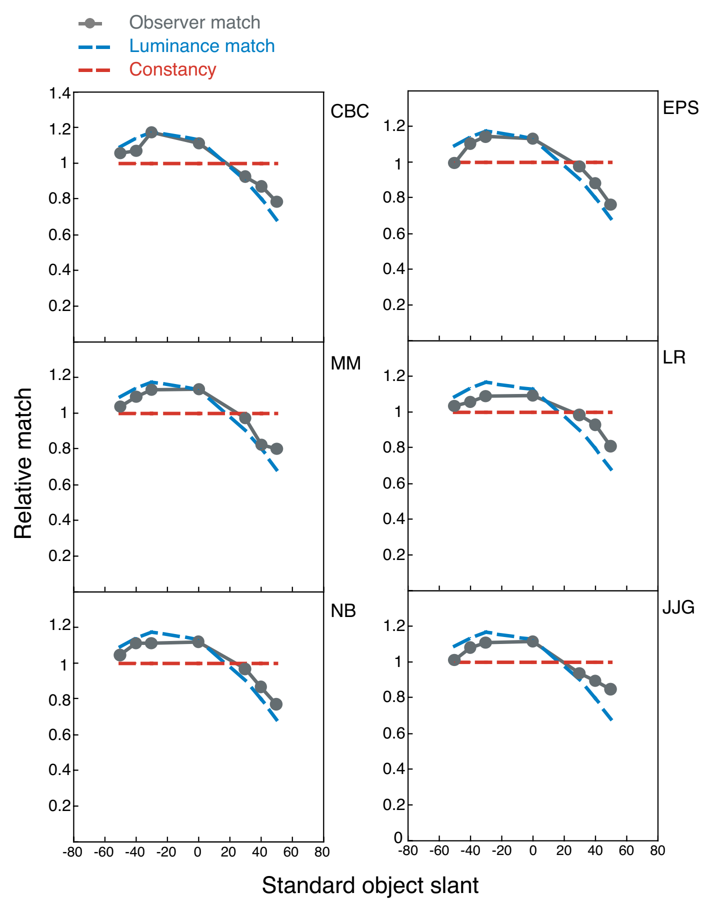

Figure 17. Data replotted from Boyaci et al. (2003) in the same format as Figure 5. Data were normalized and averaged over the two standard reflectances they used in the same manner as we aggregated our data. Luminance matching and lightness constancy predictions were scaled to the data using our normalization procedures. The figure was produced from tabulated data kindly provided by H. Boyaci.

further, perhaps by employing a wider range of instructions or by developing other methods of distinguishing perceptual and strategic aspects of lightness matching.

Some indication that observers explicitly employed different strategies is provided by written answers they gave to questions administered after they had completed the experiment. In response to the question, "Can you describe in words what aspects of the stimuli you paid more attention to when performing the matching task?," observer IBO from Experiment 2 wrote, ". . .I tried to keep in mind that the orientation of the card will change the color." This observer's data show considerable lightness constancy. Observer FGB from Experiment 3, whose data are on the luminance matching end of the spectrum, responded to the same question with "Not sure – just matched the color." Our overall impression is that there is some correlation between observers' responses as to what they were doing and their measured performance, and that this correlation is stronger than that between report and actual instructions given. It is difficult to quantify these impressions. An important issue for future research is how to assess and manipulate the strategies used by observers to perform lightness matching in complex scenes.

Other aspects of observers' reports

The evaluation forms of Experiment 3 also included questions that asked about the light sources and standard objects used in the experiment. Only 3 of 14 observers reported the correct number and location of the two light sources used. Six out of the remaining 11 observers reported that only one light source position was used but did indicate that some aspect of the lighting (presumably intensity) varied from trial to trial. Only 2 of 14 observers correctly reported the number of standard object reflectances used in the experiment, which indicates that observers were probably not explicitly memorizing standard object reflectances.

Use of real illuminants and objects

Our experiments were conduced with real illuminated objects, viewed binocularly. This is a feature they share with the early experimental work (Hochberg & Beck, 1954; Flock & Freedberg, 1970; Epstein, 1961; Gilchrist, 1980; see also Bloj et al., 1999). More recently, the availability of sophisticated computer graphics programs has made it possible to render synthetic three-dimensional scenes, and stimuli created in this way have been used to study the interaction of geometry and lightness (Knill & Kersten, 1991; Pessoa et al., 1996; Boyaci et al., 2003). Although synthetically generated stimuli offer important advantages in terms of the range of stimulus manipulations that can easily be implemented, it remains an open question as to how performance measured with synthetic stimuli relates to performance measured with actual objects.

We expect that over the next several years, this question will receive increasing attention. We believe that the class of experiments reported here, which combine the use of real stimuli and the type of parametric manipulations that often motivate the use of synthetic stimuli, will play an important role in improving our understanding of the relation between results obtained with synthetic stimuli and performance for more natural viewing.

Implications for models

Our data have a number of implications for models of lightness perception in complex three-dimensional scenes. First, the individual differences we measured require that any quantitative model contain parameters that can account for these differences. A model that simply predicts lightness as a function of the stimulus will not be able to account satisfactorily for our data.

The classic account of lightness constancy emphasizes the role of contrast coding, that is the luminance relation between a test region and its local surround (e.g., Wallach, 1948) or some other reference region in the scene (e.g., Land & McCann, 1971; see Brainard & Wandell, 1986). Earlier studies showing an interaction between perceived geometry and lightness challenge the completeness of such accounts, because in these studies the perceived lightness is often manipulated without any change to the luminance relations in the scene (Mach, 1886/1959; Hochberg & Beck, 1954; Gilchrist, 1980; Gilchrist, Delman, & Jacobsen, 1983). Our results add to this challenge – a contrast account of our data would predict a close approximation to luminance matching, because the available references in the image change little if at all with the slant of the test. It is also not clear how to model individual differences within a contrast coding approach.

Gilchrist et al. (1999; see also, Adelson, 1999) advocate a theory of lightness perception (*anchoring theory*) in which geometry acts through its role in segmenting the scene. To apply these ideas to our experiment, one would conjecture that changing the slant of the standard card affects the set of objects in the scene with which it is grouped. For any given grouping, the perceived lightness of the standard card would be determined primarily by the relative luminances of the objects within the group, which Gilchrist et al. (1999) would call the standard card's "framework." It does seem clear that grouping-like effects can have a large effect on the role of geometry in perceived lightness (e.g., Gilchrist, 1980), and our experiments do not speak directly to that question. On the other hand, because our standard card is presented in isolation, it is not obvious how to determine which other objects in the scene it might be grouped within an anchoring account. Indeed, we do not find that the grouping rules of anchoring theory are currently sufficiently well specified to allow quantitative (or even qualitative) prediction of our data. In addition, for anchoring theory to account quantitatively for our data, the grouping rules would have to be specified in a manner that allowed them to vary between observers.

To develop a model that can account quantitatively for individual variability requires a model containing a parametric description of the effect of slant on perceived lightness. Brainard and colleagues (Speigle & Brainard, 1996; Brainard, Brunt, & Speigle, 1997) have suggested that one way to parameterize the form of the transformation induced by context is through an equivalent illuminant model. In such models, the observer is assumed to be correctly performing a constancy computation, with the one exception that their estimate of the illuminant deviates from the actual illuminant. Thus the observer's estimate of the illuminant parameterizes the transformation induced by context. Boyaci et al. (2003) have successfully employed an equivalent illuminant model to account for their lightness matches as a function of slant. In the companion study (Bloj et al., 2004), we formulate such a model and evaluate how well it can account for our data.

Appendix

This "Appendix" describes the procedure we used to aggregate the matches $r_{i,\theta}$ over standard reflectance to obtain the normalized relative matches $\bar{r}_\theta^{norm}$.

For each observer, we fit the matches $r_{i,\theta}$ with a function of the form

$$r_{i,\theta} = A\frac{(k_\theta s_i)^n}{(k_\theta s_i)^n + 1}, \quad (3)$$

where s_i is the reflectance of the i^{th} standard, k_θ is a constant that depends on θ, and A and n are constants that are independent of θ. This form has two important features. First, it is separable in slant θ and standard reflectance s_i, so that the dependence of match reflectance on s_i is the same for each θ up to multiplicative constant k_θ. Second, it can account for a fairly wide range of monotonically increasing forms of the dependence of the $r_{i,\theta}$ on reflectance s_i when θ is held fixed.

Equation 3 was fit to each observer's data using numerical search. The search procedure found the values of parameters A, n, and k_θ that minimized the error ε defined by

$$\varepsilon = \sum_{i,\theta,j}\frac{\left[r_{i,\theta,j} - A\frac{(k_\theta s_i)^n}{(k_\theta s_i)^n + 1}\right]^2}{\hat{\sigma}_{i,\theta}^2}. \quad (4)$$

In this expression, the index j indicates replications of the measurement and the expression $\hat{\sigma}_{i,\theta}^2$ is an estimate of the measurement variance of matches $r_{i,\theta,j}$. The estimate $\hat{\sigma}_{i,\theta}^2$ was obtained from the mean match $r_{i,\theta}$ under the assumption of Weber's Law behavior (see Figure 4). In fitting, any matches $r_{i,\theta,j}$ that were at the minimum or maximum of the palette, or where the observer indicated that a satisfactory match was not possible, were excluded from the calculation of the fit error ε.

Given the fit of Equation 3, normalized relative matches $\bar{r}_\theta^{norm}$ were obtained through

$$\bar{r}_\theta^{norm} = k_\theta / \bar{k}, \quad (5)$$

where $\bar{k}$ is the average over θ of k_θ. The choice of $\bar{r}_\theta^{norm}$ as proportional to k_θ may be understood as follows. Suppose that the form of the function relating the $r_{i,\theta}$ to standard object reflectances s_i were linear for each θ: $r_{i,\theta} = k_\theta s_i$. In this case, it is clear that for any fixed standard object reflectance, the matches $r_{i,\theta}$ considered as a function of slant are proportional k_θ. Equation 5 preserves this interpretation in the face of a nonlinearity between perceived lightness and palette match.

The normalization by $\bar{k}$ in Equation 5 simply makes the mean of the $\bar{r}_\theta^{norm}$ over θ equal to 1. Normalization of the data seem sensible given that we have aggregated across standard reflectance, and also because the palette provides a somewhat arbitrary lightness standard. For example, one might expect a shift in the actual matches had the palette samples been surrounded by a highly reflective surface rather than the low-reflectance surround we employed. This shift would not be of interest here, where the focus is on the dependence of lightness on standard slant.

Acknowledgments

This work was supported by National Institutes of Health Grant EY 10016. John Andrews-Labenski, Mike Suplick, and Leigh Checcio helped design and construct the apparatus. We thank B. Backus, J. Hillis, A. Gilchrist, L. Maloney, H. Boyaci, J. Nachmias, and S. Sternberg for useful discussions and comments. H. Boyaci provided us with the data from Boyaci et. al (2003) in tabular format.

Commercialrelationships: none.
Corresponding author: David H. Brainard.
Email: brainard@psych.upenn.edu.
Address: Department of Psychology, University of Pennsylvania, Suite 302C, 3401Walnut Street, Philadelphia, PA 19104.

Footnotes

[1]The reflectance of each of the 36 paint mixtures was obtained by comparing the light reflected from each card to that reflected from a reflectance standard (PhotoResearch RS-2). Measurements were between 380 and 780 nm at 4-nm steps (PhotoResearch PR-650), and reflectance was obtained by averaging computed reflectance over the visual spectrum.

[2]Each datum shown in Figure 3 is the mean of matches set in three separate sessions. To resample each datum, we averaged a random draw of three samples (with replacement) from the corresponding three matches. We then ag-

gregated the resampled data using the same procedure that we used to generate the normalized relative matches shown in Figure 5. This resampling and aggregation were then repeated 1000 times and the variation in the resampled data was used to find the confidence intervals.

References

Adelson, E. H. (1999). Lightness perception and lightness illusions. In M. Gazzaniga (Ed.), *The new cognitive neurosciences* (2nd ed., pp. 339-351). Cambridge, MA: MIT Press.

Arend, L. E., & Goldstein, R. (1987). Simultaneous constancy, lightness, and brightness. *Journal of the Optical Society of America A*, *4*, 2281-2286. [PubMed]

Arend, L. E., & Reeves, A. (1986). Simultaneous color constancy. *Journal of the Optical Society of America A, 3*, 1743-1751. [PubMed]

Bauml, K. H. (1999). Simultaneous color constancy: How surface color perception varies with the illuminant. *Vision Research, 39*, 1531-1550. [PubMed]

Bloj, A., Kersten, D., & Hurlbert, A. C. (1999). Perception of three-dimensional shape influences colour perception through mutual illumination. *Nature, 402*, 877-879. [PubMed]

Bloj, M., Ripamonti, C., Mitha, K., Greenwald, S., Hauck, R., & Brainard, D. H. (2004). An equivalent illuminant model for the effect of surface slant on perceived lightness. *Journal of Vision*, *4*(9), 735-746, http://journalofvision.org/4/9/6/, doi:10.1167/4.9.6. [PubMed][Article]

Bloj, M. G., & Hurlbert, A. C. (2002). An empirical study of the traditional Mach card effect. *Perception, 31*, 233-246. [PubMed]

Boyaci, H., Maloney, L. T., & Hersh, S. (2003). The effect of perceived surface orientation on perceived surface albedo in binocularly viewed scenes. *Journal of Vision, 3*(8), 541-553, http://journalofvision.org/3/8/2/, doi:10.1167/3.8.2. [PubMed][Article]

Brainard, D. H., Brunt, W. A., & Speigle, J. M. (1997). Color constancy in the nearly natural image. 1. Asymmetric matches. *Journal of the Optical Society of America A, 14*, 2091-2110. [PubMed]

Brainard, D. H., & Wandell, B. A. (1986). Analysis of the retinex theory of color vision. *Journal of the Optical Society of America A, 3*, 1651-1661. [PubMed]

Epstein, W. (1961). Phenomenal orientation and perceived achromatic color. *Journal of Psychology, 52*, 51-53.

Flock, H. R., & Freedberg, E. (1970). Perceived angle of incidence and achromatic surface color. *Perception and Psychophysics, 8*, 251-256.

Gilchrist, A., Kossyfidis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X., Spehar, B., Annan, V., & Economou, E. (1999). An anchoring theory of lightness perception. *Psychological Review, 106*, 795-834. [PubMed]

Gilchrist, A. L. (1980). When does perceived lightness depend on perceived spatial arrangement? *Perception and Psychophysics, 28*, 527-538. [PubMed]

Gilchrist, A. L., Delman, S., & Jacobsen, A. (1983). The classification and integration of edges as critical to the perception of reflectance and illumination. *Perception and Psychophysics, 33*, 425-436. [PubMed]

Henneman, R. (1935). A photometric study of the perception of object color. *Archives of Psychology*, *179*, 5-89.

Hochberg, J. E., & Beck, J. (1954). Apparent spatial arrangement and perceived brightness. *Journal of Experimental Psychology, 47*, 263-266. [PubMed]

Knill, D. C., & Kersten, D. (1991). Apparent surface curvature affects lightness perception. *Nature, 351*, 228-230. [PubMed]

Land, E. H., & McCann, J. J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America*, *61*, 1-11. [PubMed]

Mach, E. (1959). *The Analysis of Sensations*. New York: Dover. (Original work published 1886)

Pessoa, L., Mingolla, E., & Arend, L. E. (1996). The perception of lightness in 3-D curved objects. *Perception and Psychophysics, 58*(8), 1293-1305. [PubMed]

Speigle, J. M., & Brainard, D. H. (1996). Luminosity thresholds: Effects of test chromaticity and ambient illumination. *Journal of the Optical Society of America A, 13*(3), 436-451. [PubMed]

Van Ee, R., & Erkelens, C. J. (1995). Temporal aspects of binocular slant perception. *Vision Research, 36*(1), 43-51. [PubMed]

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). An empirical explanation of brightness. *Proceedings of the National Academy of Sciences U.S.A., 95*, 13301-13306. [PubMed][Article]