Perception, 1999, volume 28, pages 217 – 242

Estimator reliability and distance scaling in stereoscopic slant perception

Benjamin T Backus, Martin S Banks¶§
Vision Science Program and School of Optometry (¶ also Department of Psychology), University of California, Berkeley, CA 94720-2020, USA; e-mail: marty@john.berkeley.edu
Paper submitted in conjunction with the Applied Vision Association Conference/Workshop on Depth Perception, Guildford, Surrey, UK, 4 September 1997; revised version received 7 September 1998

Abstract. When a horizontal or vertical magnifier is placed before one eye, a frontoparallel surface appears slanted. It appears slanted away from the eye with horizontal magnification (geometric effect) and toward the eye with vertical magnification (induced effect). According to current theory, the apparent slant in the geometric and induced effects should increase with viewing distance. The geometric effect does scale with distance, but there are conflicting reports as to whether the induced effect does. Ogle (1938 *Archives of Ophthalmology* **20** 604 – 623) reported that settings in slant-nulling tasks increase systematically with viewing distance, but Gillam et al (1988 *Perception & Psychophysics* **44** 473 – 483) and Rogers et al (1995 *Perception* **24** Supplement, 33) reported that settings in slant-estimation tasks do not. We re-examined this apparent contradiction. First, we conducted two experiments whose results are consistent with the literature and thus replicate the apparent contradiction. Next, we analyzed the signals available for stereoscopic slant perception and developed a general model of perceived slant. The model is based on the assumption that the visual system knows the reliability of various slant-estimation methods for the viewing situation under consideration. The model's behavior explains the contradiction in the literature. The model also predicts that, by manipulating eye position, apparent slant can be made to increase with distance for vertical, but not for horizontal, magnification. This prediction was confirmed experimentally.

1 Introduction

The visual system uses sensor measurements to construct representations of the environment. These measurements are subject to error, so the resulting representations are never exactly correct. If more than one method is available for constructing a representation, the results of these methods will generally differ. Combining the results (by, for example, averaging) usually yields a more accurate final representation. The means by which the visual system combines results from different information sources can be examined by using the cue-conflict paradigm in which information sources, or signals, are manipulated independently.

Here we take up the surprisingly complicated question of how a change in viewing distance affects the apparent slant of a stereoscopically defined surface. We start by describing an apparent contradiction in the literature: in slant-nulling tasks, distance has a pronounced effect on perceived slant, but in slant-estimation tasks it does not. We report two experiments that replicate the apparent contradiction. Next, we develop a general model for predicting apparent slant. The model is based on the assumption that the visual system knows the reliability of various methods for estimating slant as a function of the viewing situation. We present calculations that show how the actual reliabilities change from one situation to another, and show how this can explain the contradiction. Finally, we show that a remarkable prediction of the model is confirmed experimentally: apparent slant from vertical magnification can be made to scale with distance if eye position is manipulated.

§ Author to whom all correspondence and requests for reprints should be addressed.

1.1 *Some necessary definitions*

Figure 1 schematizes the geometry of the binocular viewing situation and defines some needed terms. The left panel is an overhead (plan) view of the situation. A surface is fixated binocularly. The lines of sight from the two eyes are represented by the dashed lines; the cyclopean line of sight is the diagonal solid line. The distance d to the fixation point is measured along the cyclopean line of sight. The slant S is the angle between the surface and the objective gaze-normal plane (a plane perpendicular to the cyclopean line of sight). The angles γ and μ are the eyes' version and vergence, respectively. Positive slant (S) and azimuth (γ) are defined counterclockwise viewed from above.

The right panel of figure 1 is an oblique view from behind the observer. The fixation point on the surface patch is indicated by the small white circle. Angles α_L and α_R are the horizontal angles (in the visual plane) subtended by the patch at the left and right eyes, respectively; the horizontal size ratio (HSR) is defined as α_L/α_R (Rogers and Bradshaw 1993). Angles β_L and β_R are the vertical angles subtended by the surface patch; the vertical size ratio (VSR) is β_L/β_R. A vertical magnifier of strength m is shown in front of the left eye; it causes the proximal VSR to become $m\beta_L/\beta_R$. 'Objective slant' will refer to the equivalent physical slant of a simulated surface prior to magnification. For example, a frontoparallel surface fixated in the head's median plane has an objective slant of 0°; the proximal HSR, and thus apparent slant, would vary with the strength of a horizontal magnifier.

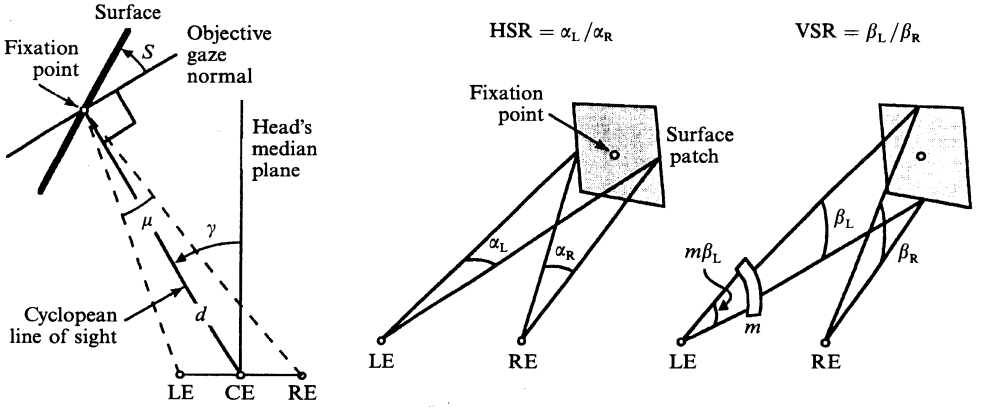

Figure 1. Binocular viewing geometry. The left panel is an overhead (plan) view of the situation considered here. LE and RE refer to the left and right eyes. CE is the cyclopean eye, positioned at the midpoint between the left and right eyes. The head's median plane passes through the cyclopean eye and is perpendicular to the interocular axis. A surface is fixated by the two eyes at the fixation point. The lines of sight from the two eyes are represented by the dashed lines and the cyclopean line of sight by the diagonal solid line. Distance d to the fixation point is measured along the cyclopean line of sight. Slant S is the angle between the surface and the objective gaze-normal plane (a plane perpendicular to the cyclopean line of sight). The angles γ and μ are the eyes' version and vergence, respectively. S, γ, and μ are angles in the visual plane, which is defined by the eyes and the fixation point (ie the plane of the figure). Positive slant (S) and azimuth (γ) are defined counterclockwise viewed from above. The right panel is an oblique view from behind the observer. The fixation point on the surface patch is indicated by the small white circle. Angles α_L and α_R are the horizontal angles subtended by the patch at the left and right eyes. The horizontal size ratio (HSR) is α_L/α_R. Angles β_L and β_R are the vertical angles subtended by the surface patch. The vertical size ratio (VSR) is β_L/β_R. A vertical magnifier of strength m is shown in front of the left eye; it causes the proximal VSR to become $m\beta_L/\beta_R$. [Adapted from Backus et al (1999); copyright Pergamon Press.]

1.2 *The geometric and induced effects*

In the geometric effect, a horizontal magnifier placed before one eye causes a frontoparallel surface to appear slanted about a vertical axis; if the right eye's image is magnified, the right side of the surface appears farther than the left. In the induced effect, a vertical magnifier placed before one eye also causes a frontoparallel surface to appear slanted, but in this case the side nearer to the magnifier looks closer. The geometric and induced effects are predicted by the relationship between slant S and the signals HSR, VSR, and μ, as defined in figure 1.[1] This relationship is:

$$S = S_{\mathrm{HSR,VSR}} \approx -\arctan\left(\frac{1}{\mu}\ln\frac{\mathrm{HSR}}{\mathrm{VSR}}\right) \tag{1}$$

(see Backus et al 1999, for derivation).

To see how the geometric effect is predicted from equation (1), consider a frontoparallel surface fixated in the head's median plane. Before magnification, HSR = VSR = 1. Placing a horizontal magnifier before the left eye magnifies the HSR term by factor m, so equation (1) becomes:

$$S_{\mathrm{apparent}} \approx -\arctan\left(\frac{\ln m}{\mu}\right). \tag{2}$$

For $m > 1$, $S < 0$, so the objectively frontoparallel surface will appear slanted in a clockwise direction (right side near).

The induced effect is also predicted. Again HSR = VSR = 1 for a frontoparallel surface (at the head's median plane). Placing a vertical magnifier before the left eye magnifies the VSR term by factor m, so equation (1) becomes:

$$S_{\mathrm{apparent}} \approx \arctan\left(\frac{\ln m}{\mu}\right). \tag{3}$$

For $m > 1$, $S > 0$, so the objectively frontoparallel surface appears slanted in an anticlockwise direction (left side near).

According to equations (2) and (3), apparent slant in the geometric and induced effects should be equal in magnitude and opposite in sign for a given magnification. In addition, apparent slant should increase equally with viewing distance because, as distance increases, μ decreases, so the magnitude of S will increase. This prediction can be visualized in figures 2 and 3. The figures are overhead views of the visual plane. The X and Z axes represent lateral and forward position with respect to the midpoint between the eyes. The thick black segments represent surface patches. The gray curves represent isovergence (iso-μ) circles. In figure 2, the surface patches straddle the head's median plane, so VSR = 1. The patches all have slants such that HSR = 1.02. As distance increases, the slant magnitude increases. In figure 3, the two black circles represent iso-VSR contours, VSR = 1.02 and 1/1.04 on the left and right, respectively. For a surface to yield HSR = 1, it must be tangent to an isovergence circle. Thus, as distance increases, the slant magnitude again increases.[2] This is the basis for the prediction that the apparent slant of a truly frontoparallel surface, seen with a vertical magnifier before one eye, should increase with distance.

[1] The signal μ can be measured in two ways: from an extraretinal, eye-position signal or from the horizontal and vertical disparities (Mayhew and Longuet-Higgins 1982; Gillam and Lawergren 1983; Rogers and Bradshaw 1993, 1995; Gårding et al 1995; Backus et al 1999). A variety of other cues, such as accommodation, could in principle provide information about μ.

[2] It can also be seen that beyond a distance of about 150 cm, there is no point in space, and hence no realizable surface, for which VSR = 1/1.04. Technically, therefore, the more general prediction is that apparent slant should increase as μ decreases, not as distance increases.

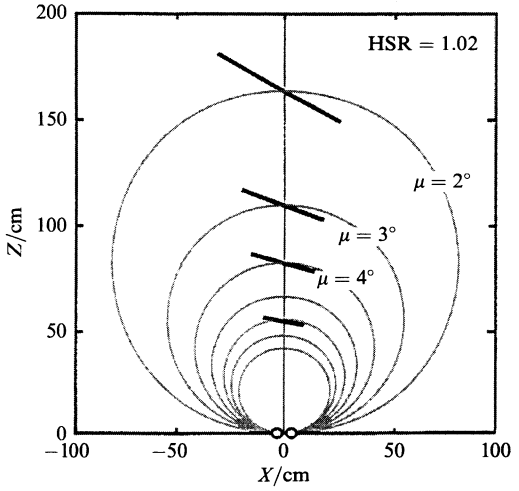

Figure 2. The geometric effect and distance. The figure is an overhead view through the visual plane. The X and Z axes represent lateral and forward position, respectively, with respect to the cyclopean eye which is at the origin. The small circles near the origin represent the eyes. The gray circular contours are isovergence contours; from the smallest circle to the largest, $\mu = 8°$, 7°, 6°, 5°, 4°, 3°, and 2°. The black line segments represent surface patches for which the horizontal size ratio (HSR) = 1.02. The vertical size ratio (VSR) = 1 because the patches lie in the head's median plane. Notice that the surface slant associated with HSR = 1.02 increases with increasing distance. When 2% horizontal magnification is applied to the left eye and the observer views an objectively gaze-normal surface, HSR = 1.02, so the magnitude of the perceived slant ought to increase with distance as indicated by the black line segments.

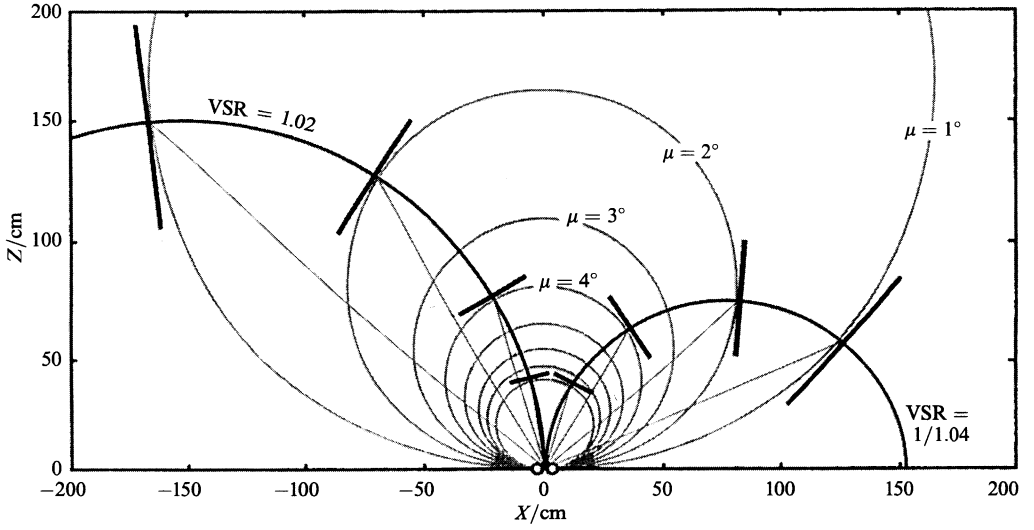

Figure 3. The induced effect and distance. The figure is an overhead view through the visual plane. The black contours show locations for which the vertical size ratio (VSR) is constant; the contour on the left represents VSR = 1.02 and the one on the right represents VSR = 1/1.04. The gray circular contours are isovergence contours. The black line segments to the left of the head's median plane represent surfaces for which HSR = 1 and VSR = 1.02. The black lines to the right represent surfaces for which HSR = 1 and VSR = 1/1.04. Notice that the surface slant associated with those values of VSR and HSR increases with increasing distance from the cyclopean eye. When 2% vertical magnification is applied to the left eye and the observer is viewing an objectively gaze-normal surface, VSR = 1.02 and HSR = 1, so the perceived slant should increase with distance as indicated by the black line segments to the left of the median plane. A similar argument applies to 4% vertical magnification applied to the right eye.

We call estimation of slant from HSR, VSR, and μ [equation (1)] *slant estimation by HSR and VSR*. Two other methods of slant estimation are generally available in experimental measurements of the geometric and induced effects. First, it can be shown that:

$$S = S_{\mathrm{HSR,EP}} \approx -\arctan\left(\frac{1}{\mu}\ln \mathrm{HSR} - \tan\gamma\right). \tag{4}$$

Thus, the visual system can estimate slant from HSR, μ, and γ. We assume that γ (the eyes' version) is measured from extraretinal eye-position signals. Equation (4) thus describes *slant estimation by HSR and eye position* (Banks and Backus 1998).[3] Second, if there are residual nonstereo cues, such as perspective, they allow *slant estimation by nonstereo cues*.[4]

Each of these methods of estimating slant defines a *slant estimator*. We call them $\hat{S}_{\mathrm{HSR,VSR}}$, $\hat{S}_{\mathrm{HSR,EP}}$, and $\hat{S}_{\mathrm{nonstereo}}$, respectively. Use of $\hat{S}_{\mathrm{HSR,EP}}$ alone would predict the geometric effect (because HSR is affected by a horizontal magnifier) but not the induced effect [because none of the signals in equation (4) is affected by a vertical magnifier]. Use of $\hat{S}_{\mathrm{nonstereo}}$ alone would predict neither effect (Banks and Backus 1998).

The geometric and induced effects have been measured as a function of viewing distance by Gillam et al (1988), Rogers et al (1995), Ogle (1938), and Amigo (1972). Gillam et al and Rogers et al used a slant-estimation procedure which is portrayed in figure 4. In this task, observers are shown an objectively gaze-normal surface and are asked to indicate the perceived slant. Data from one observer in Gillam et al (1988) are

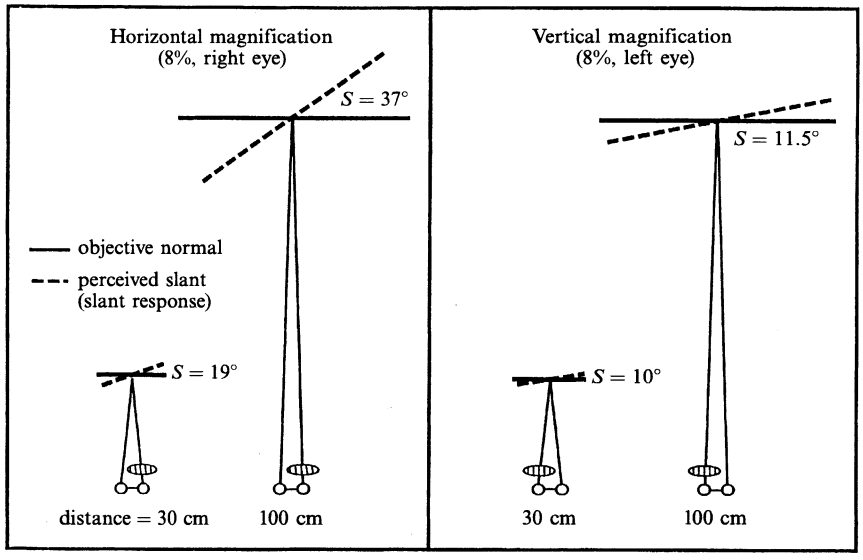

Figure 4. Schematic of the slant-estimation task employed by Gillam et al (1988). Observers viewed an objectively gaze-normal surface (solid lines) and indicated its perceived slant (dashed lines). The left panel shows results for two viewing distances when 8% horizontal magnification was applied to the right eye. At 30 and 100 cm, observer BG indicated perceived slants of 19° and 37°, respectively. The right panel shows results when 8% vertical magnification was applied to the left eye (the data were actually collected with magnification applied to the right eye; we changed the eye and slant direction to make the figure). At 30 and 100 cm, BG indicated perceived slants of 10° and 11.5°, respectively. Thus, perceived slant increased much more with increasing viewing distance in the geometric effect than in the induced effect.

[3] This approximation is improved from that given in Banks and Backus (1998). See Backus et al (1999) for derivation.

[4] Gillam et al (1988) noted that an additional cue based on motion may account for their failure to replicate Ogle's (1938) slant-nulling results.

presented in figure 5. With horizontal magnification (left panel), perceived slant increased with viewing distance as predicted by equation (2). Rogers et al (1995) replicated this observation. With vertical magnification (right panel), Gillam and colleagues reported that perceived slant did not increase with distance. Thus, the predictions exemplified by equation (3) were not confirmed. Again Rogers and colleagues replicated this finding.

Ogle (1938) used a slant-nulling procedure which is portrayed in figure 6. In this task, observers are presented a surface and are asked to manipulate its slant until it

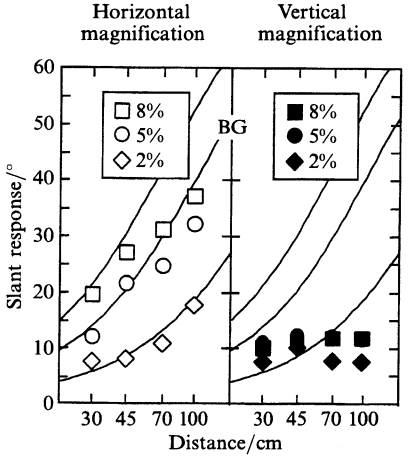

Figure 5. Results from the slant-estimation experiment of Gillam et al (1988). The responses of observer BG are plotted in absolute degrees as a function of viewing distance. The left and right panels show data for horizontal and vertical magnification, respectively. Values are shown for magnifications of 2%, 5%, and 8%. Notice that slant responses increased systematically with distance with horizontal magnification, but did not vary with distance with vertical magnification. The curves represent predicted settings based on the slant estimated from HSR and VSR [equation (1)].

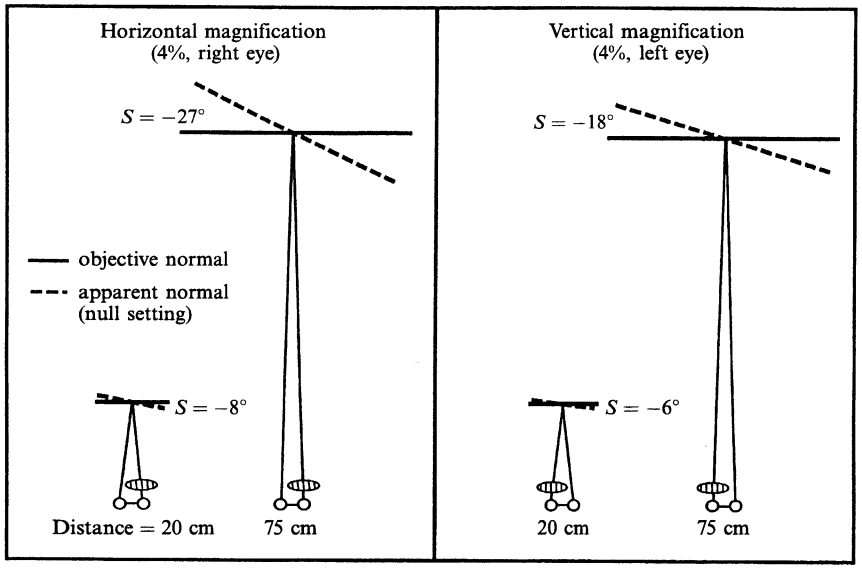

Figure 6. Schematic of the slant-nulling task employed by Ogle (1938). Observers adjusted the slant of a surface (dashed lines) until it appeared to be gaze normal (solid lines). The left panel shows two conditions when 4% horizontal magnification was applied to the right eye. At viewing distances of 20 and 75 cm, observer RHD rotated the surface to slants of −8° and −27°, respectively, to make it appear gaze normal. The right panel shows two conditions when 4% vertical magnification was applied to the left eye. At distances of 20 and 75 cm, RHD rotated the surface to −6° and −18°, respectively. (Separate data were actually collected for both types of magnification in each eye; we averaged the null settings after reversing the sign of the slant for one eye's magnification.) Thus, the surface slant required to create the appearance of a gaze-normal surface increased significantly with increasing viewing distance in both the geometric effect and induced effect.

appears gaze normal. The nulling task differs from the estimation task because, after perceived slant has been nulled in the induced effect, HSR no longer equals 1. However, according to equation (1), the plane should appear gaze normal when HSR = VSR. For vertical magnification (VSR = m), the objective slant that gives rise to HSR = m is:

$$S_{\text{null}} \approx -\arctan\left(\frac{\ln m}{\mu}\right). \tag{5}$$

The equation for horizontal magnification is the same except for a sign reversal. Thus, equation (1) predicts responses of equal size to a given magnification for the slant-estimation task [equation (3)] and the slant-nulling task [equation (5)].

Data from one of Ogle's observers are presented in figure 7. As predicted by equation (5), the magnitude of the null setting increased with viewing distance for both horizontal (left panel) and vertical (right panel) magnification, although the distance effect was smaller for vertical magnification.

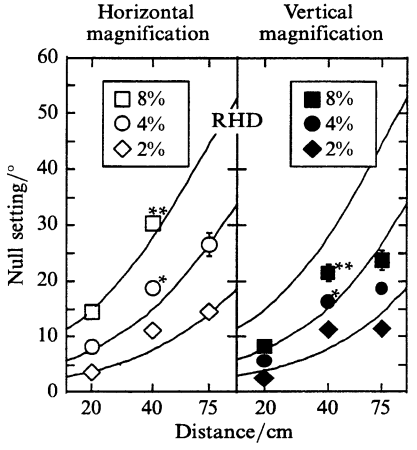

Figure 7. Results from the slant-nulling experiment of Ogle (1938). The slants of the nulling surface are plotted in absolute degrees as a function of viewing distance for observer RHD. The left and right panels show data for horizontal and vertical magnification, respectively, of 2%, 4%, and 8% (data from the two eyes are averaged). (At 40 cm, the magnifications were actually 2%, 4.5%, and 8.2%; asterisks indicate the affected data points.) Notice that slant responses increased systematically with distance for both horizontal and vertical magnification. The curves represent predicted settings based on the slant estimated from HSR and VSR [equation (1)].

In another nulling experiment, Amigo (1972) used sparse stimuli and reported a large geometric effect that scaled with distance and a small induced effect that did not. We will consider Amigo's results again in section 7.

Thus, there are two puzzles. Why in slant-estimation experiments does the geometric effect vary predictably with viewing distance while the induced effect does not (Gillam et al 1988; Rogers et al 1995)? Why is this pattern not observed in slant-nulling tasks where the geometric and induced effects both vary predictably with distance (Ogle 1938)?

2 Experiment 1: Geometric and induced effects with slant estimation

The differences in the findings of the two sets of studies could be artifactual, rather than being a consequence of the difference in task. For this reason, we attempted to replicate the previous work. We began by using a slant-estimation procedure to measure geometric and induced effects at a variety of viewing distances.

2.1 *Methods*

There were five observers: the authors (BTB and MSB) and three adults who were unaware of the experimental hypotheses (SJF, WHF, TLY). All had normal binocular vision. BTB, WHF, and TLY wore contact lenses to correct their refractive errors.

Stimuli were displayed on a haploscope consisting of two large monochrome CRT displays each seen in a mirror by one eye (Banks and Backus 1998; Backus et al 1999). The mirror and CRT were attached to an armature that rotated about a vertical axis passing through the eye's center of rotation. The face of each CRT was always

perpendicular to the line of sight from the eye to the center of the screen. A small target on each screen served as the fixation aid. Head position was fixed with a bite bar. The room was completely dark, and a black aperture occluded the frames of the monitors. Only the white dots within the display were visible.

Despite the short viewing distance of 40 cm, the visual locations of the dots in our displays were specified to within $\sim$ 30 s of arc. This high level of spatial precision was achieved by use of two procedures: anti-aliasing and spatial calibration (Backus et al 1999).

The stimuli consisted of sparse random-dot displays. To create the displays, 300 dots were placed on a simulated vertical plane located 19.1, 57.3, or 172 cm in front of the midpoint of the interocular axis. The dots were randomly distributed within a circle subtending 25 deg at the cyclopean eye. Because the angular subtense of the display was approximately constant, the outline shape of the display gave rise to a nonstereo cue indicating zero slant. Dot size and intensity also did not vary with slant or magnification. Before magnification, the images were geometrically correct for a frontoparallel plane in forward gaze at the specified viewing distance. Then, the left or right image was magnified either vertically or horizontally by 1.25%, 2.5%, 5%, or 7.5%. The displays preserved size disparities in the Hess coordinate system. Specifically, in order to preserve size disparities in the direction orthogonal to the magnification, vertical magnification changed the latitudes but not the azimuths of image points; horizontal magnification changed the azimuths but not the latitudes. Use of another coordinate system, such as the Helmholtz or Fick system, would have made little difference in how magnification was applied to the images [Schor et al (1994) describe these coordinate systems].

The displays contained a number of cues to the plane's distance. Vergence (μ) and the pattern of vertical disparities ($\partial \mathrm{VSR}/\partial\gamma$) specified the simulated distance (19.1, 57.3, or 172 cm). Accommodation specified the actual distance to the screens (40 cm). Dot brightness and size and overall display size did not specify a distance, but did indicate that no change in distance occurred with a change in simulated distance. The arms of the haploscope were fixed in position throughout experiments 1 and 2, so the lateral position of the stimuli on the CRTs varied with distance.

Observers initiated stimulus presentations with a key press. The fixation marker appeared at the center of the display and observers were instructed to fixate it whenever it was present. The stimulus plane appeared 1 s later and remained visible (along with the fixation marker) for 3 s. A blank screen then appeared for 250 ms, followed by a response display that consisted of two lines visible only to the left eye. Observers adjusted the angle between the lines until it matched the perceived slant of the stimulus plane (van Ee and Erkelens 1996). Observers were not given feedback. There were six or seven trials at each combination of distance and magnification.

2.2 *Results*

Figure 8 shows predictions and results for horizontal (left panels) and vertical (right panels) magnification. Slant response magnitudes are plotted as a function of viewing distance. The predictions from the three means of slant estimation are displayed in the upper left panels. Use of HSR and VSR [equation (1)] yields the four functions represented by solid curves, one for each magnification value. Use of HSR and eye position [equation (4)] yields the same curves in the case of horizontal magnification, but yields a horizontal line at 0° in the case of vertical magnification. Although nonstereo cues to slant were minimized, they were not eliminated completely; use of these cues yields horizontal lines at 0° for both magnifications.

The data are plotted in separate panels for each of the five observers: horizontal and vertical magnification data are shown in the left and right panels, respectively. In each panel, average slant responses are plotted as a function of distance. There were no

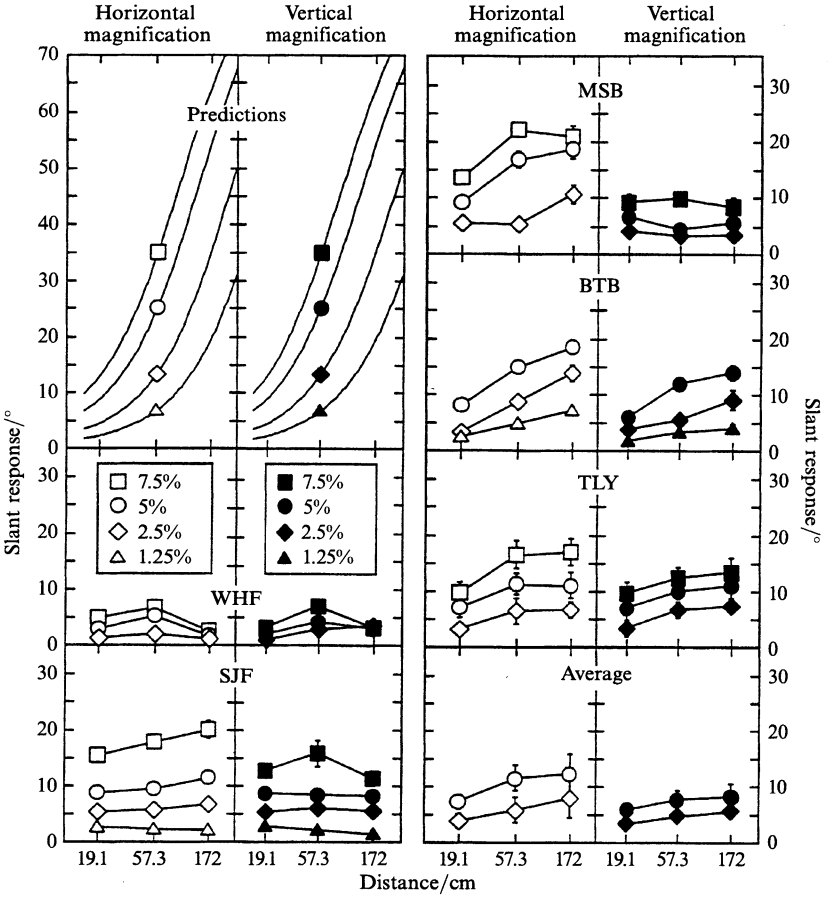

Figure 8. Predictions and results for experiment 1. Predictions for single slant estimators are shown in the upper left pair of panels. Predicted slant responses are plotted as a function of distance. The left and right panels display predictions for horizontal and vertical magnification, respectively. Values are shown for magnifications of 1.25%, 2.5%, 5%, and 7.5%. The curves in the horizontal magnification panel represent the predictions for the estimators $\hat{S}_{HSR,EP}$ and $\hat{S}_{HSR,VSR}$; the estimator $\hat{S}_{nonstereo}$ predicts a horizontal line at 0°. The curves in the vertical magnification panel represent the predictions for $\hat{S}_{HSR,EP}$; $\hat{S}_{HSR,EP}$ and $\hat{S}_{nonstereo}$ predict a horizontal line at 0°. The results are shown in the other panels, a separate pair for each observer. Average slant response is plotted as a function of distance; slant responses obtained with magnification in the left eye have been averaged with sign-inverted responses obtained with magnification in the right eye. Error bars represent ±1 standard error, with standard error calculated as $[(\sigma_L^2 + \sigma_R^2)/(2N)]^{1/2}$, $N = 6$ or 7, where σ_L and σ_R are the standard deviations for the left-eye and right-eye magnification data, respectively (twelve or fourteen trials total per data point). The average data are shown in the lower right pair of panels; error bars are ±1 standard error, computed across observers.

systematic differences in the data obtained with left-eye and right-eye magnifications, so the slant responses were averaged after reversing the sign of the response for one eye.

These data show large individual differences in the amount of slant seen for a given proximal stimulus. For horizontal magnification (left panels), four observers (SJF, MSB, BTB, TLY) exhibited larger slant responses at greater distances (the effect in SJF was small at low magnifications). Their data are consistent with the literature, although the increase in slant was smaller than predicted from equations (1) and (4) and smaller than reported by Gillam et al (1988). One observer (WHF) did not exhibit a systematic effect of viewing distance. For vertical magnification (right panels), three observers (WHF, SJF, MSB) exhibited little increase in slant response with increasing distance.

Their data replicate the reports by Gillam et al (1988) and Rogers et al (1995) that the magnitude of the induced effect does not increase with viewing distance. However, two observers (BTB, TLY) exhibited a systematic increase in slant response with increasing distance which does not replicate the earlier reports.[5]

The results of experiment 1 are thus generally consistent with the observations of Gillam et al (1988) and Rogers et al (1995). With horizontal magnification (geometric effect), the perceived slant of an objectively gaze-normal surface increased as the distance increased; this statement is valid for four of the five observers and is quite consistent with the average data. With vertical magnification (induced effect), perceived slant did not increase systematically with increasing distance. Again, we observed differences among our five observers. The four observers who showed systematic effects of distance all showed a greater effect for horizontal than for vertical magnification.

3 Experiment 2: Slant nulling in forward gaze

Using a slant-nulling procedure (figure 6), Ogle (1938) observed systematic increases in the magnitude of the induced effect when stimulus distance was increased. We next asked whether we could replicate his observations with the same stimuli and observers that were used in experiment 1.

3.1 *Methods*

There were four observers: the authors (BTB and MSB) and two adults who were unaware of the experimental hypotheses (TLY and WHF).

The equipment, stimuli, and experimental procedure were identical to those employed in experiment 1 except that observers now adjusted the slant of the stimulus plane until it appeared normal to the cyclopean line of sight. They did so by making key presses to indicate the direction in which the plane appeared rotated from gaze normal. The stimulus was visible until the key press, at which point it was extinguished and then redrawn at another slant. Observers made a series of adjustments until they were satisfied that the stimulus appeared gaze normal. They were not given feedback.

3.2 *Results*

Figure 9 shows predictions and results for horizontal (left panels) and vertical (right panels) magnification. The predictions from the two stereoscopic methods of slant estimation are shown in each panel. Use of HSR and VSR [equation (1)] yields the three functions represented by solid curves, one for each magnification value. Use of HSR and eye position [equation (4)] yields the same curves in the case of horizontal magnification, but yields a horizontal line at 0° in the case of vertical magnification. Nonstereo cues always indicated a slant of 0° and, therefore, could not influence the slant setting (Banks

[5] A two-way ANOVA was performed for each observer at each magnification level. The factors were Magnification Type (horizontal and vertical) and Distance. There were 12 or 14 replicates per cell. The main effect of Distance was highly significant ($p < 0.001$) for BTB and TLY at all three magnifications, for WHF at 5% and 7.5% magnification, and for MSB at 5% magnification. It was very significant ($p < 0.01$) for MSB at 2.5% and 7.5%, and significant ($p < 0.05$) for WHF at 2.5% magnification. It was not significant for SJF at any magnification. The main effect of Magnification Type was highly significant for MSB and BTB at all magnifications, and for SJF and TLY at 7.5%. It was very significant for SJF at 5%. It was not significant elsewhere. The interaction between Distance and Magnification Type was very significant for MSB at all magnifications; it was significant for WHF and BTB at 2.5%, and for SJF and TLY at 7.5%. From figure 8, it is clear that all significant main effects and interactions are in the expected direction, except in the case of WHF, who showed on average a decrease in slant response at the largest distance, and a larger slant response for 2.5% vertical than for 2.5% horizontal magnification at 172 cm. In particular, the data as a whole show that responses are greater for horizontal than for vertical magnification, and that the size of this difference grows with increasing distance.

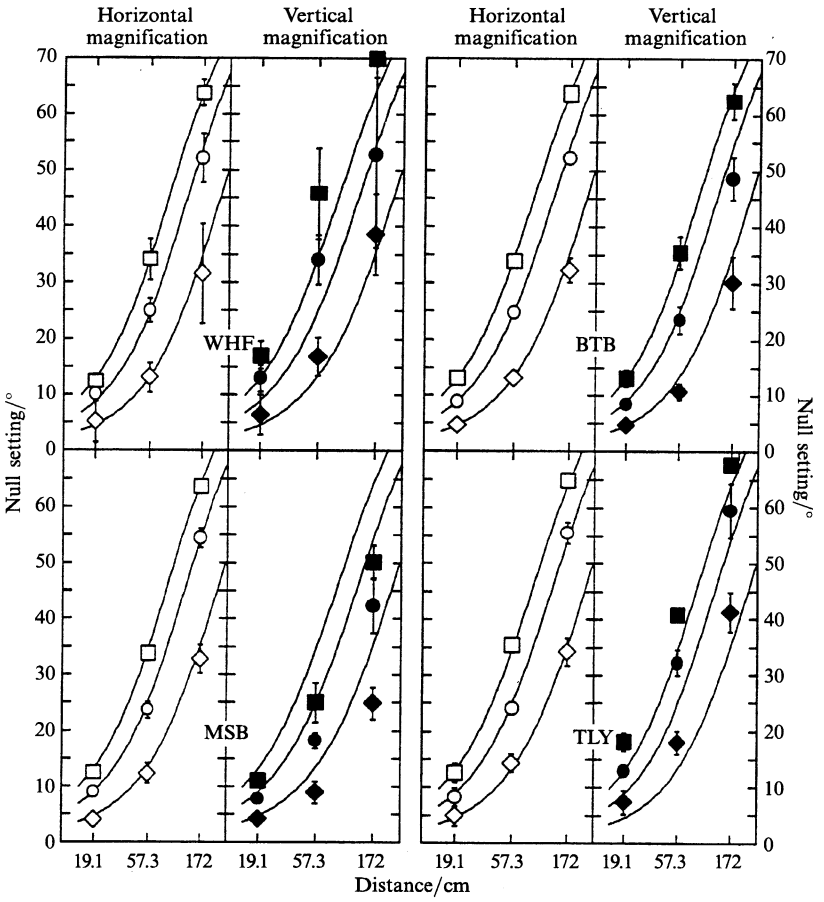

Figure 9. Predictions and results for experiment 2. Predicted slant-nulling settings for single slant estimators are represented in each panel by the curves; they are plotted as a function of distance. The left and right panels display predictions and results for horizontal and vertical magnification, respectively. Values are shown for magnifications of 2.5% (diamonds), 5% (circles), and 7.5% (squares). For horizontal magnification, the curves represent the predictions for the estimators $\hat{S}_{HSR,EP}$ and $\hat{S}_{HSR,VSR}$; the estimator $\hat{S}_{nonstereo}$ predicts a horizontal line at 0°. For vertical magnification, the curves represent the predictions for $\hat{S}_{HSR,VSR}$; $\hat{S}_{HSR,EP}$ and $\hat{S}_{nonstereo}$ predict a horizontal line at 0°. The results are shown in the other panels, a separate pair for each observer. In each panel, the average null setting is plotted as a function of distance; setting magnitudes obtained with magnification applied to the left eye have been averaged with setting magnitudes obtained with magnification applied to the right eye. Error bars represent ±1 standard deviation (average standard deviation for left-eye and right-eye data). Each data point represents ten trials.

and Backus 1998); this feature differs from the experiment of Ogle (1938) in which a real surface was used and gave rise to nonstereo cues consistent with its true slant.

The data are plotted in separate panels for each of the four observers: horizontal and vertical magnification data in the left and right panels, respectively. In each panel, average slant-null settings are plotted as a function of distance. There were no systematic differences in the data obtained with left-eye and right-eye magnifications, so the settings were averaged after reversing sign for one eye.

These data show clear and systematic effects of viewing distance for both horizontal and vertical magnification. The effect of distance when one eye's image is magnified horizontally is shown in the left panels. As expected, all observers showed systematic increases in null settings with increasing distance. The settings were very close to the predicted values. The effect of distance when one eye's image is magnified vertically is

shown in the right panels. All four observers exhibited systematic increases in null settings with increasing distance; this replicates Ogle's report that the magnitude of the induced effect increases with viewing distance. In addition, the magnitudes of the settings were quite similar to the predicted values for observers BTB, TLY, and WHF; they fell somewhat lower than the predicted values for observer MSB.

The results of experiment 2 are quite consistent with Ogle's observations. With horizontal and vertical magnification, the slant of an apparently gaze-normal surface increased as the distance to the surface became greater.

Even though the observations of Gillam et al (1988) and Rogers et al (1995) seem to contradict those of Ogle (1938), we basically replicated both sets of observations. Thus, the puzzle remains: why in a slant-nulling task does the induced effect scale with viewing distance while in a slant-estimation task it does not?

4 Signal combination and slant estimation: Theory

The predictions derived for experiments 1 and 2 are for single slant estimators [eg equations (1) and (4)]. To understand why distance scaling occurs in some circumstances and not others, we need to consider how those estimator outputs are combined to form a final slant estimate. To do so, we will now examine how noise in signal measurement affects the reliability of each slant estimate, and how the estimates could be combined to yield the most reliable final slant estimate. We have already described three slant estimators—slant estimation by HSR and VSR ($\hat{S}_{\mathrm{HSR,\,VSR}}$), slant estimation by HSR and eye position ($\hat{S}_{\mathrm{HSR,\,EP}}$), and slant estimation by nonstereo cues ($\hat{S}_{\mathrm{nonstereo}}$)—each of which can in principle yield a slant percept in the viewing situations associated with induced-effect and geometric-effect measurements. Each estimator's output is a function of the environmental signals it uses. In the geometric effect, the only signal altered by the magnification is HSR, so we expect $\hat{S}_{\mathrm{HSR,\,VSR}}$ and $\hat{S}_{\mathrm{HSR,\,EP}}$ to be affected and to agree with one another; $\hat{S}_{\mathrm{nonstereo}}$ is unaffected by the magnification and will, therefore, disagree with the outputs of the other two estimators. In the induced effect, only VSR is altered, so we expect $\hat{S}_{\mathrm{HSR,\,VSR}}$ to disagree with $\hat{S}_{\mathrm{HSR,\,EP}}$ and $\hat{S}_{\mathrm{nonstereo}}$ which agree with one another (Banks and Backus 1998). Given the conflicts among the estimator outputs, how should the visual system derive a final slant estimate?

Figure 10 shows our model which is an example of a *modified weak fusion* model (Landy et al 1995). In weak fusion models (Clark and Yuille 1990), the reliability of each signal is estimated from ancillary measures, which include cues that are not used in the estimate itself. Examples of ancillary cues might include the local texture content within a region of the scene or the distance to an object. Estimates based on the separate signals are combined in a weighted average; the weight of each estimator is proportional to its reciprocal variance.

The signals HSR, VSR, μ, γ, and the texture gradient are environmental signals that have exact values. We call the visual system's measurements of the first four signals $\widehat{\mathrm{HSR}}$, $\widehat{\mathrm{VSR}}$, $\hat{\mu}$, and $\hat{\gamma}$; they are, of course, subject to error. The texture gradient also contributes to a slant estimate (namely $\hat{S}_{\mathrm{nonstereo}}$), but we do not discuss this or other nonstereo signals in sufficient detail to label them separately.

After making some assumptions about the noises in the signal measurements, we can calculate each slant estimator's *reliability*. We define an estimator's reliability as the reciprocal of its variance. If X and Y are two estimators with reliabilities r_X and r_Y, the estimator with highest reliability is

$$Z = \frac{r_X}{r_X + r_Y} X + \frac{r_Y}{r_X + r_Y} Y$$

(Blake et al 1993; Landy et al 1995). If X and Y are independent, then Z has reliability $r_Z = r_X + r_Y$. If estimators use a common signal, they will not be independent, and the

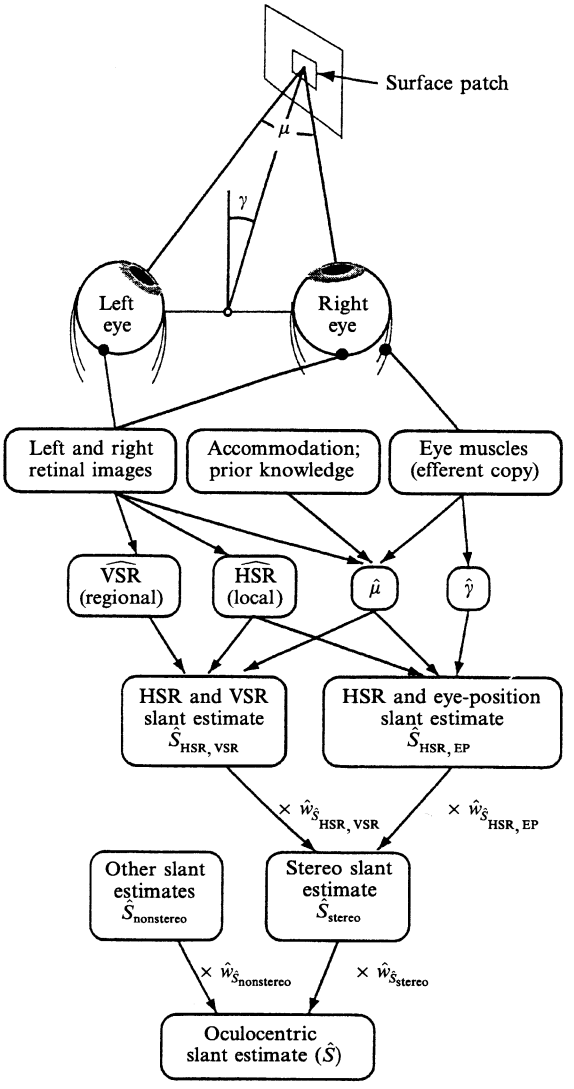

Figure 10. Model of stereoscopic slant estimation. The visual system measures $\widehat{\text{HSR}}$ and $\widehat{\text{VSR}}$ from the retinal images, and $\hat{\gamma}$ and $\hat{\mu}$ from the eye muscles ($\hat{\mu}$ can also be measured from the retinal images). From those measurements, it derives a slant estimate ($\hat{S}_{\text{HSR, VSR}}$) based on HSR and VSR and an estimate ($\hat{S}_{\text{HSR, EP}}$) based on HSR and eye position. These estimates are added, after being weighted by their normalized reliabilities $\left(\hat{w}_{\hat{S}_{\text{HSR, VSR}}}\right.$ and $\left.\hat{w}_{\hat{S}_{\text{HSR, EP}}}\right)$, to yield a stereoscopic slant estimate ($\hat{S}_{\text{stereo}}$). The visual system also measures signals based on nonstereoscopic cues ($\hat{S}_{\text{nonstereo}}$). The final slant estimate ($\hat{S}$) is the sum of $\hat{S}_{\text{stereo}}$ and $\hat{S}_{\text{nonstereo}}$ after they have been weighted by their normalized reliabilities ($\hat{w}_{\hat{S}_{\text{stereo}}}$ and $\hat{w}_{\hat{S}_{\text{nonstereo}}}$). Other cues to distance, such as accommodation, might affect performance.

reliability of the weighted average will usually be less than the sum of the estimators' separate reliabilities; this is the case for $\hat{S}_{\text{HSR, VSR}}$ and $\hat{S}_{\text{HSR, EP}}$ because they both use $\widehat{\text{HSR}}$ and $\hat{\mu}$.

What should the visual system do if the relevant slant estimators report high reliabilities for their estimates, but the estimator outputs disagree? In many cases, the disagreement is caused by a failure of one of the estimators and, in that event, a linearly weighted average no longer provides the best estimate of the parameter; a robust procedure is needed to downweight the errant estimate (Huber 1981; Landy et al 1995).

The plateau observed at high magnifications in the induced effect is an example of this phenomenon (Banks and Backus 1998). A complete theory of slant estimation would include a robust procedure, but a linear weighting scheme suffices for understanding how distance affects apparent slant in the geometric and induced effects. Examples of data from cue-conflict experiments that are consistent with a linear scheme are provided by Buell and Hafter (1991), Johnston et al (1993), Rogers and Bradshaw (1995), Bradshaw et al (1996), Heller and Trahiotis (1996), and Backus et al (1999).[6] That the visual (and auditory) systems sometimes behave simply and consistently in the face of radically conflicting information would seem to be an important discovery.

To summarize, signal noise causes error in the three slant estimates. The reliability of each estimator ought to vary as a function of the viewing situation. To obtain the most reliable final estimate possible, the visual system should weight each estimate according to its reliability in the current viewing situation. If we assume that the visual system does this, we can model the final slant percept as:

$$\hat{S} = \frac{r_{\text{stereo}}}{r_{\text{stereo}} + r_{\text{nonstereo}}} \hat{S}_{\text{stereo}} + \frac{r_{\text{nonstereo}}}{r_{\text{stereo}} + r_{\text{nonstereo}}} \hat{S}_{\text{nonstereo}} , \tag{6}$$

where

$$\hat{S}_{\text{stereo}} = \frac{r_{\text{HSR, VSR}}}{r_{\text{HSR, VSR}} + r_{\text{HSR, EP}}} \hat{S}_{\text{HSR, VSR}} + \frac{r_{\text{HSR, EP}}}{r_{\text{HSR, VSR}} + r_{\text{HSR, EP}}} \hat{S}_{\text{HSR, EP}} . \tag{7}$$

We next determine what the reliabilities are, given some assumptions about the noises in the signals.

5 Signal combination and slant estimation: Simulations

5.1 *Method*

To simulate slant estimation for experiments 1 and 2, we rewrote equations (1) and (4) with each signal replaced by signal plus noise. This allows error in the slant estimate to be written as a Taylor-series function of the signal errors, expanded about the true signal values. This analysis revealed complicated interactions between the true signal values and the amounts of noise in the signals, so it was necessary to use Monte Carlo simulations to determine the reliability of each estimator. The noise distribution associated with each signal measurement is not directly measurable, so we assumed Gaussian distributions with means of 0 and made plausible assumptions about the standard deviations in the measurements (see Appendix). According to the viewing situation, we drew a sample from each signal distribution, then calculated the three slant estimates from the samples. This was repeated 40 000 times. The reliability of each estimate (calculated from its distribution) was used to weight the estimates, according to equations (7) and (6), respectively.

5.2 *Results*

We present the results in a series of four figures. Figures 11 and 12 show the reliabilities of four slant estimators, $\hat{S}_{\text{HSR, VSR}}$, $\hat{S}_{\text{HSR, EP}}$, $\hat{S}_{\text{stereo}}$, and $\hat{S}_{\text{nonstereo}}$, as a function of surface distance and slant. Figure 13 shows how the final slant estimate was computed. Figure 14 summarizes the simulation results from figure 13 and shows how changes in the gain of the measurement of vergence affect the outcome.

Figure 11 shows the reliabilities of the stereoscopic slant estimators, $\hat{S}_{\text{HSR, VSR}}$ and $\hat{S}_{\text{HSR, EP}}$, for three viewing situations. From top to bottom, the situations are natural viewing, viewing with 5% horizontal magnification in the left eye, and viewing with 5% vertical magnification in the left eye. The eyes are fixating a surface in the head's

[6] Studies of stereoscopic shape perception have certainly documented nonlinear interactions between competing cues. Some earlier examples include Wallach et al (1963), Dosher et al (1986), Bülthoff and Mallot (1988), Stevens and Brookes (1988), and Buckley and Frisby (1993).

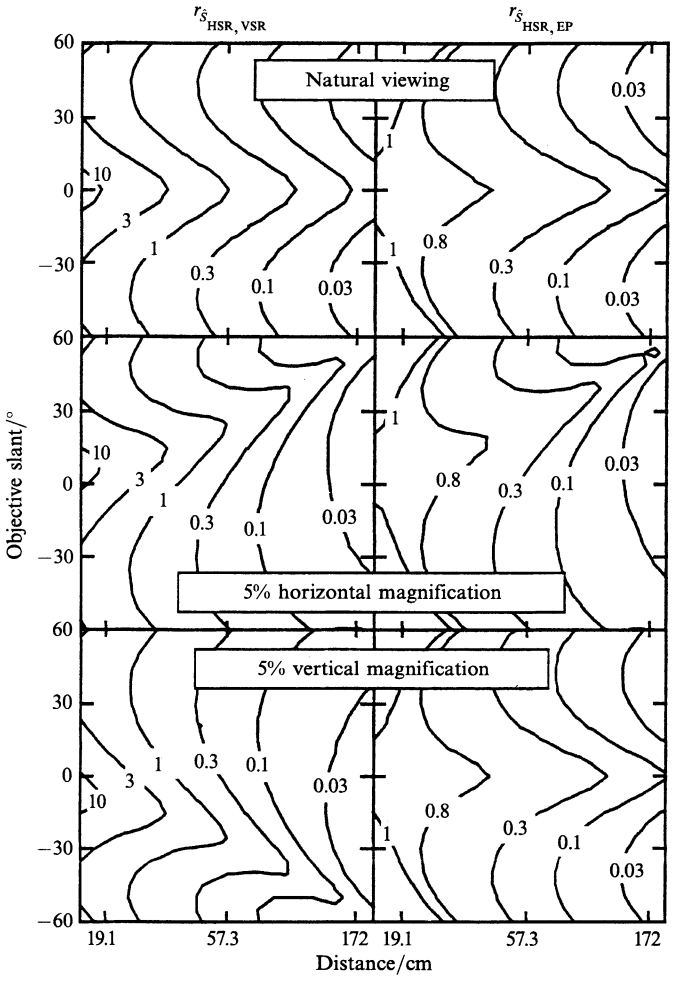

Figure 11. The reliabilities of stereoscopic slant estimators, $\hat{S}_{HSR,VSR}$ and $\hat{S}_{HSR,EP}$. The eyes are in forward gaze ($\gamma = 0°$) fixating a surface whose slant varies about a vertical axis. The abscissae are the distances to the surface and the ordinates are the objective slants. From top to bottom, the three pairs of panels represent natural viewing, viewing with 5% horizontal magnification in the left eye, and viewing with 5% vertical magnification in the left eye. The panels on the left show the reliability of $\hat{S}_{HSR,VSR}$ as a function of distance and slant, and the panels on the right show the reliability of $\hat{S}_{HSR,EP}$. The contours in each panel are isoreliability contours; each is labeled with its actual reliability. Each panel was generated from a grid of 25 distances by 25 slants. Reliability at each point in the grid was the reciprocal variance of 40 000 slant estimates obtained in the Monte Carlo simulation.

median plane, so $\gamma = 0°$. The abscissae in figure 11 are distances to the surface and the ordinates are objective slants of the surface. The left panels show the reliability (the reciprocal of the computed variance) of $\hat{S}_{HSR,VSR}$ and the right panels the reliability of $\hat{S}_{HSR,EP}$. Isoreliability contours are shown.

For natural viewing, $\hat{S}_{HSR,VSR}$ (left panel) and $\hat{S}_{HSR,EP}$ (right panel) are most reliable when the objective slant is 0°. As distance increases, the reliabilities of $\hat{S}_{HSR,VSR}$ and $\hat{S}_{HSR,EP}$ decrease. The reliability of $\hat{S}_{HSR,EP}$ is limited at small distances (see section 7).

For horizontal magnification (middle panels), the most reliable estimated slant occurs at objective slants that increase with distance, resulting in the ridge of reliability for $\hat{S}_{HSR,VSR}$ and $\hat{S}_{HSR,EP}$ that runs from the middle of the panel on the left to the upper right corner.

For vertical magnification (bottom panels), the most reliable estimated slant occurs for $\hat{S}_{HSR,VSR}$ at objective slants that decrease with distance, resulting in the ridge of

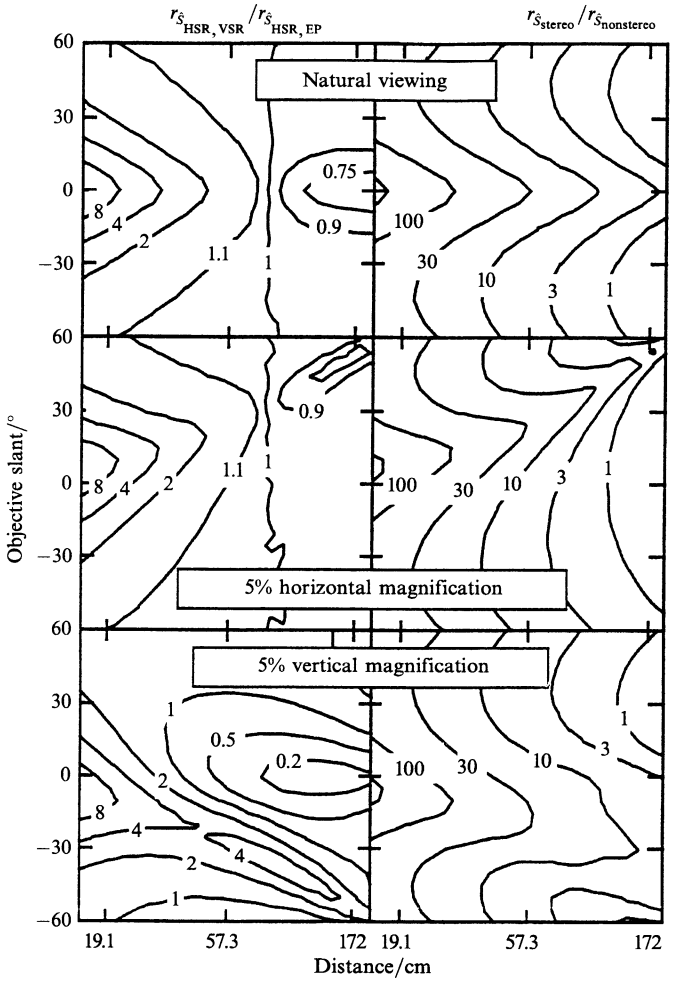

Figure 12. The relative reliabilities of slant estimators. The eyes are in forward gaze fixating a surface whose slant varies about a vertical axis. The abscissae are the surface distances and the ordinates are the objective slants. From top to bottom, the pairs of panels represent natural viewing, viewing with 5% horizontal magnification in the left eye, and viewing with 5% vertical magnification in the left eye. The left panels show the relative reliabilities of the stereoscopic estimators: reliability of $\hat{S}_{HSR,VSR}$ divided by reliability of $\hat{S}_{HSR,EP}$. The right panels show the relative reliabilities of the stereo and nonstereo estimators: reliability of $\hat{S}_{stereo}$ divided by reliability of $\hat{S}_{nonstereo}$. The contours represent constant reliability ratios. The reliability of $\hat{S}_{nonstereo}$ was assumed to be constant at 0.04 deg^{-2} for all distances and objective slants. Each panel was generated from a 25 by 25 grid. Reliability at each point in the grid was the reciprocal variance of 40 000 slant estimates obtained in the Monte Carlo simulation.

reliability that runs from middle to lower right. It occurs for $\hat{S}_{HSR,EP}$ at zero slant (same as natural viewing).

Figure 12 shows the relative reliabilities of $\hat{S}_{HSR,VSR}$ and $\hat{S}_{HSR,EP}$ (left panels) and $\hat{S}_{stereo}$ and $\hat{S}_{nonstereo}$ (right panels) for the viewing conditions of figure 11. In the left panel, a ratio of 1 implies equal weighting for $\hat{S}_{HSR,VSR}$ and $\hat{S}_{HSR,EP}$ in equation (7); in the right panels, it implies equal weighting for $\hat{S}_{stereo}$ and $\hat{S}_{nonstereo}$ in equation (6). The results for vertical magnification are of greatest interest. $\hat{S}_{HSR,VSR}$ is generally more reliable than $\hat{S}_{HSR,EP}$ at short viewing distances. However, when objective slant is 0°, as in the slant-estimation task, and viewing distance increases, $\hat{S}_{HSR,EP}$ becomes increasingly reliable relative to $\hat{S}_{HSR,VSR}$. As a consequence, the slant estimate of 0° from $\hat{S}_{HSR,EP}$ is given

more and more weight at the same time that the slant estimate from $\hat{S}_{HSR,VSR}$ is increasing. The two trade off, so $\hat{S}_{stereo}$ should not scale with distance in an estimation task. What about the slant-nulling task? $\hat{S}_{HSR,VSR}$ reports that slant has been nulled when HSR = VSR. The long ridge in the bottom left panel of figure 12 corresponds to slants for which HSR = VSR; along this ridge, $\hat{S}_{HSR,VSR}$ is roughly 4 times more reliable than $\hat{S}_{HSR,EP}$ over a wide range of distances. In the nulling task, the observer rotates the stimulus away from objective normal such that the surface slant approaches this ridge. Then the stereoscopic estimate in the nulling task is determined primarily by $\hat{S}_{HSR,VSR}$, which forces the objective slant to increase with distance.(7)

Figure 13 shows the model's slant estimates for the conditions of experiments 1 and 2. The average slant estimates obtained from four different estimators ($\hat{S}_{HSR,VSR}$, $\hat{S}_{HSR,EP}$, $\hat{S}_{nonstereo}$, and $\hat{S}$) are plotted against proximal HSR within the stimulus. Different panels

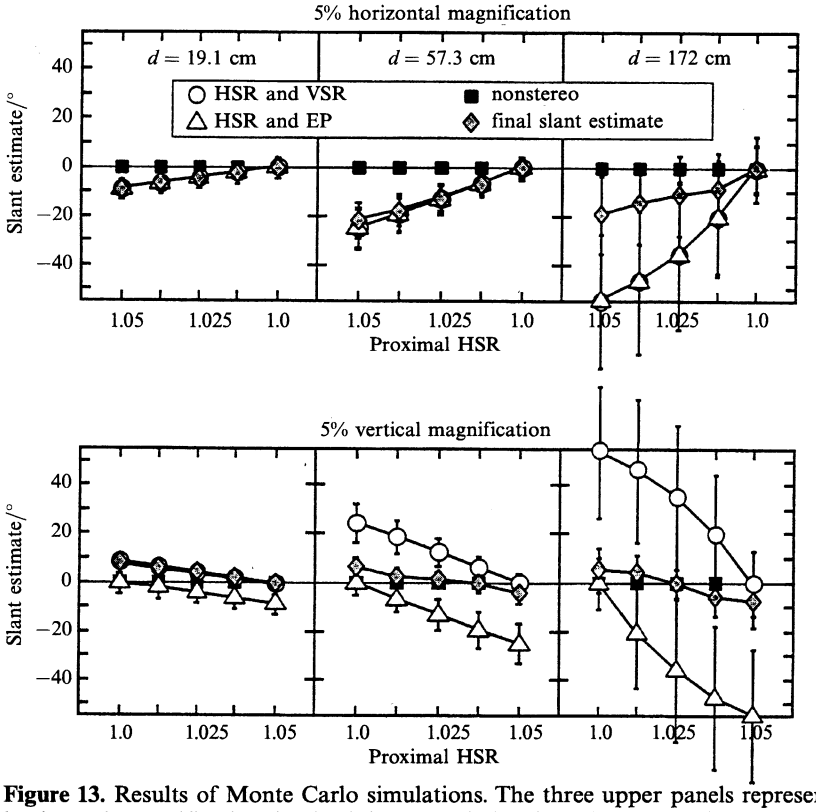

Figure 13. Results of Monte Carlo simulations. The three upper panels represent viewing with 5% horizontal magnification in the left eye and the three lower panels represent viewing with 5% vertical magnification in the left eye. The abscissae are the proximal HSR (HSR at the retinae created by the objective slant of the surface plus any horizontal magnification). The ordinates are the average estimates from slant estimators: $\hat{S}_{HSR,VSR}$ (circles), $\hat{S}_{HSR,EP}$ (triangles), $\hat{S}_{nonstereo}$ (squares), and $\hat{S}$ (diamonds). From left to right, the panels represent simulation results for viewing distances of 19.1, 57.3, and 172 cm. Error bars represent standard deviations for each estimator, magnified fourfold. The standard deviation of $\hat{S}_{nonstereo}$ was held constant at 5°.

(7) The lower right panel of figure 12 shows a large region of high reliability for $\hat{S}_{stereo}$ at far distances when objective slant is negative. This is a consequence of negative correlation between the two stereoscopic slant estimates: $\hat{S}_{HSR,VSR}$ is positive and $\hat{S}_{HSR,EP}$ is negative, so that noise in $\hat{\mu}$ on a given trial has opposite effects on the two estimates [see equations (1) and (4)]. We have no reason to suppose that the visual system would actually be able to measure, or base its weighting upon, the reduced variance of $\hat{S}_{stereo}$ that occurs in this special cue-conflict situation. The qualitative predictions of the model are similar if one assumes, for example, that the reliability of $\hat{S}_{stereo}$ equals the maximum of the two stereo estimators' reliabilities.

show results for different distances. The leftmost diamond in each panel shows the final slant estimate in the case of the slant-estimation task (objective slant = 0°). Data at the four other values of proximal HSR in each panel show the slant estimates at other objective slants (ie as the surface is rotated to null apparent slant). The proximal HSR at which the final slant estimate equals 0° gives the model's prediction for how much rotation is needed to null apparent slant. Error bars show standard deviations, magnified fourfold. The error bars in the bottom middle panel (vertical magnification at 57.3 cm) show how, as the surface is rotated away from objective gaze normal, the reliability of $\hat{S}_{\mathrm{HSR,VSR}}$ increases relative to that of $\hat{S}_{\mathrm{HSR,EP}}$.

The top two panels of figure 14 show the final slant estimates replotted for the estimation and nulling tasks (left and right panels, respectively). For vertical magnification, distance has much less effect in the estimation than in the nulling task. For horizontal magnification, distance has somewhat less effect in the estimation task. These behaviors are similar to our experimental observations.

In the simulations above, we assumed that all signal estimators are unbiased and so the only source of error was noise. The assumption of unbiased estimation seems implausible for the eyes' vergence because there is considerable experimental evidence that human observers overestimate vergence (μ) at short distances and underestimate it at long distances (Foley 1980; Bradshaw et al 1996; Brenner and van Damme 1998). Such biases can be represented as a reduction in the gain between an environmental signal and the representation of it by the nervous system. Thus, we reran the simulations with various assumptions about the gain between the eyes' vergence (μ) and the estimate of vergence ($\hat{\mu}$). Gain_{μ} is the amplitude of the estimate divided by the amplitude of the signal; it was set to 1, 0.5, or 0.2; we assumed that $\hat{\mu} = \mu$ at the actual CRT distance of 40 cm ($\mu = 8.7°$). The results of those simulations are shown in the middle and bottom panels of figure 14. The correspondence between the human and model behaviors is closest when gain_{μ} is 0.2.

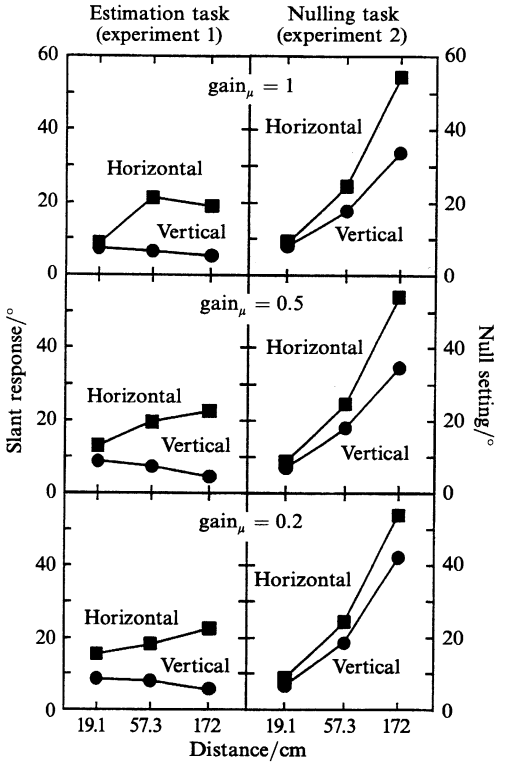

Figure 14. Simulation results for experiments 1 and 2. The left column shows simulation results for the slant-estimation task of experiment 1. Top, middle, and bottom panels display results when different values of gain_{μ} were assumed: 1, 0.5, and 0.2, respectively. In each panel, the absolute value of the model's average slant response is plotted as a function of viewing distance. The squares represent average responses for 5% horizontal magnification and the circles responses for 5% vertical magnification. The right column shows results for the slant-nulling task of experiment 2. Again, top, middle, and bottom panels show results for gain_{μ} = 1, 0.5, 0.2, respectively. In each panel, the absolute value of the average surface slant at which estimated slant ($\hat{S}$) was 0° is plotted as a function of distance.

6 Experiment 3: Slant estimation in eccentric gaze

Our simulations offer an explanation of the puzzle that motivated this work. In particular, they show that efficient use of the signals available in geometric-effect and induced-effect experiments should lead to less distance scaling in slant-estimation tasks than in slant-nulling tasks, particularly in the induced effect. A critical aspect of our explanation is that $\hat{S}_{\mathrm{HSR,EP}}$ and $\hat{S}_{\mathrm{HSR,VSR}}$ disagree when one eye's image is magnified vertically, but agree when one eye's image is magnified horizontally. We wished to determine whether this difference was critical to the failure of distance scaling in the induced effect when measured with an estimation procedure. Thus, we conducted another experiment in which we manipulated eye position in order to make $\hat{S}_{\mathrm{HSR,EP}}$ and $\hat{S}_{\mathrm{HSR,VSR}}$ agree in the induced effect and disagree in the geometric effect. We predict more distance scaling in the induced effect when extraretinal, eye-position signals are consistent with $\widehat{\mathrm{VSR}}$ because the two stereoscopic estimators, $\hat{S}_{\mathrm{HSR,EP}}$ and $\hat{S}_{\mathrm{HSR,VSR}}$, will then agree and the more reliable estimator—$\hat{S}_{\mathrm{HSR,EP}}$—will now also scale with distance. Similarly, the magnitude of $\hat{S}_{\mathrm{HSR,EP}}$ will be 0° in the geometric effect, so the geometric effect should increase less with distance than it did in forward gaze.

6.1 *Methods*

The five observers who participated in experiment 1 also participated in experiment 3. The methods were identical to those of experiment 1 with a few exceptions. First, and most importantly, the head-centric azimuth of the stimulus was varied. The azimuth was varied by rotating the haploscope arms and displacing the images on the CRTs.[8] For measurements with vertical magnification, the azimuths were either 0° or those that in natural viewing would create the VSR presented to the retinae. For measurements with horizontal magnification, the azimuths were either 0° or were set such that $\tan\gamma$ was equal to $\ln \mathrm{HSR}/\mu$; this yielded $\hat{S}_{\mathrm{HSR,EP}} = 0$ [equation (4)] and thereby created the same conflict between $\hat{S}_{\mathrm{HSR,EP}}$ and $\hat{S}_{\mathrm{HSR,VSR}}$ that occurred with vertical magnification in experiment 1. (In fact, the same set of nonzero azimuths satisfy the criteria for both vertical and horizontal magnification at a given magnification.)

Stimuli subtended 21 deg at the cyclopean eye (this small reduction in size was needed so that stimuli could be placed closer to the monitor edges). Observers indicated the perceived slant of the stimulus relative to gaze normal. Distance was 19.1 or 57.3 cm. Vertical and horizontal magnification was 5% in the left eye or 5% in the right eye. There were ten trials at each combination of distance, magnification, and gaze condition.

[8] The displacement of the images on the CRT screens yielded variation in the angle between the left and right eyes' lines of sight and normals to the CRTs. This variation created a potential slant cue, so we conducted a preliminary experiment to determine its impact. Specifically, we asked observers to view the displays monocularly and to make slant judgments. Three observers (MSB, BTB, and JMH) participated. The monitor was placed on a platform that could rotate smoothly about a vertical axis. Observers viewed the display with the left eye through an aperture that occluded the monitor frame. Viewing distance was 40 cm and the experiment was conducted in the dark. The experimenter rotated the monitor so that the face of the CRT had a slant of −12° to 12° at fixation. Dot locations were correct for a gaze-normal surface 172 cm from the left eye. Observers made sign-of-slant judgments at seven slants (random order). From these judgments, a psychometric function (percent 'right side far' vs slant) was generated. The standard deviation of the best-fitting cumulative normal was the measure of variability in slant estimates from residual nonstereo cues. The standard deviations were 4.3° (MSB), 9.8° (BTB), and 7.1° (JMH). By comparison with our modeling and experimental results, these values suggest that the greatest interactions with distance in experiment 3 will be due to conflict between $\hat{S}_{\mathrm{HSR,EP}}$ and $\hat{S}_{\mathrm{HSR,VSR}}$, not between $\hat{S}_{\mathrm{stereo}}$ and $\hat{S}_{\mathrm{nonstereo}}$. In addition, it can be shown that the direction of the interaction expected from nonstereo cues is contrary to our prediction. We conclude that nonstereo cues cannot account for the data if the data show a larger induced effect, and a smaller geometric effect, in eccentric gaze.

6.2 *Results*

Figure 15 shows the main results from experiment 3. Data when the eyes were in forward and eccentric gaze are represented by the unfilled and filled symbols, respectively. Slant responses obtained for 5% magnification in the left eye have been combined with those for 5% magnification in the right eye by averaging the magnitudes of the slant responses from those two conditions. As in experiment 1, there were clear individual differences, but the eccentric-gaze data were consistently different from the forward-gaze data (figure 8).

For vertical magnification in eccentric gaze (filled circles in right panel for each observer), every observer exhibited increasing slant responses with increasing distance. In forward gaze, slant responses were smaller and varied less systematically with distance.

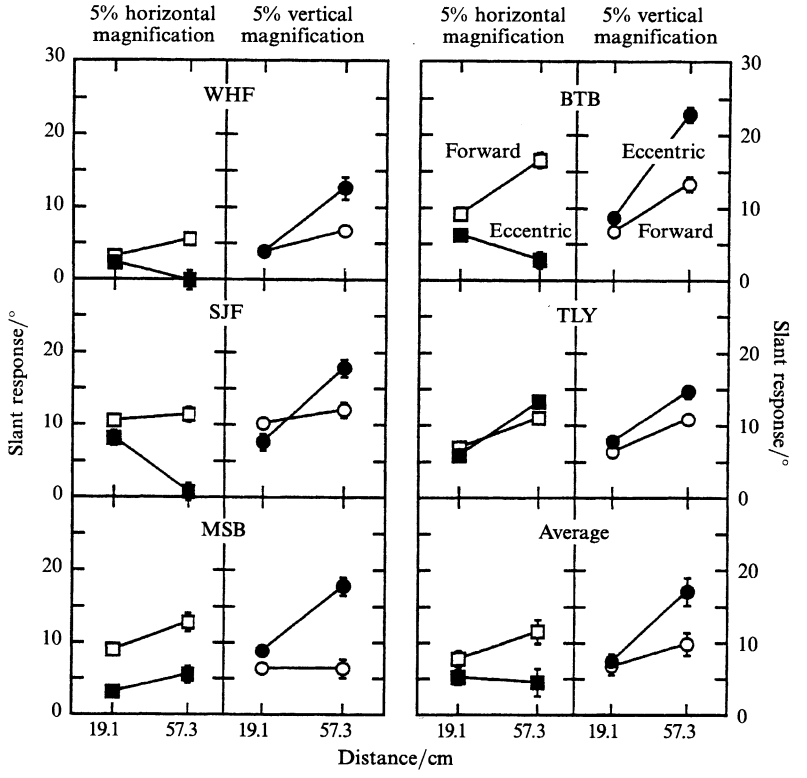

Figure 15. Results for experiment 3. Each pair of panels shows average slant responses (absolute values) for an individual observer or slant responses averaged across observers (lower right). Responses are plotted as a function of viewing distance. Left and right panels in each pair display predictions for 5% horizontal and 5% vertical magnification, respectively. Unfilled symbols represent responses when eyes were in forward gaze ($\gamma = 0°$) and filled symbols responses when eyes were in eccentric gaze. Slant response magnitudes obtained with magnification in the left eye have been averaged with response magnitudes obtained with magnification in the right eye. Error bars represent ± 1 standard error with standard error calculated as $[(\sigma_L^2 + \sigma_R^2)/(2N)]^{1/2}$, $N = 10$, where σ_L and σ_R are the standard deviations for the left-eye and right-eye magnification data, respectively (twenty trials total per data point). The average data are shown in the lower right pair of panels; error bars are ± 1 standard error, computed across observers.

For horizontal magnification in eccentric gaze (filled squares in left panels), responses were smaller and less correlated with distance than in forward gaze. Thus, as predicted, we were able to reverse the effects of distance on the induced and geometric effects by placing the eyes in eccentric gaze.

6.3 *Simulation results*

We used the model depicted in figure 10 for the conditions of experiment 3. The parameters of the model were identical to those in the simulations of experiments 1 and 2 except that the true value of γ varied in the eccentric-gaze condition. The results of the simulations are depicted in figure 16. They are quite similar to the human data, especially when gain_μ is 0.2. The effects of distance, type of magnification, and eye position are all qualitatively similar in the human and simulation results. Quantitatively, the model overestimated observers' slant responses.

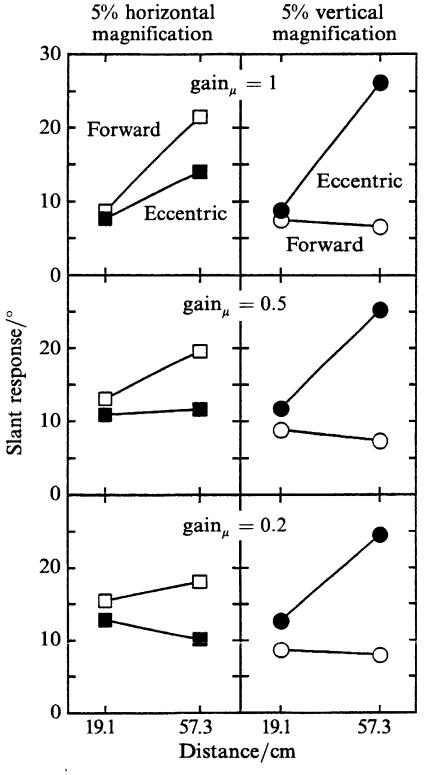

Figure 16. Simulation results for experiment 3. Left and right columns show simulation results with 5% horizontal magnification and 5% vertical magnification, respectively. Top, middle, and bottom panels display results for $\text{gain}_\mu = 1$, 0.5, 0.2, respectively. In each panel, the absolute value of the model's average slant response is plotted as a function of viewing distance. Squares represent average responses for 5% horizontal magnification and circles responses for 5% vertical magnification. Unfilled symbols represent responses with eyes in forward gaze ($\gamma = 0°$) and filled symbols responses with eyes in eccentric gaze.

7 Discussion

Our primary empirical result is confirmation of the apparently contradictory data in the literature from slant-nulling and slant-estimation experiments. Although we found noteworthy interobserver variability, we largely confirmed that distance scaling occurs in the presence of horizontal and vertical magnification when the psychophysical task is slant nulling (Ogle 1938; experiment 2) and that distance scaling occurs more for horizontal than for vertical magnification when the task is slant estimation (Gillam et al 1988; Rogers et al 1995; experiment 1). We were able to explain this apparent contradiction with an analysis of environmental signals, measurement noises, and weight assignment when combining estimator outputs.

7.1 *Comparison of modeling results and equations (1) and (4)*

We used Monte Carlo simulations to find reliabilities for the two stereoscopic slant estimators as functions of distance and objective slant (figure 11). Inspection of equations (1) and (4) yields an understanding of these results. From equation (1), it is evident that $\hat{S}_{\text{HSR, VSR}}$ will be least sensitive to error in $\hat{\mu}$ when HSR = VSR. This condition is satisfied along the ridges of high reliability in the left panels of figure 11.

From equation (4), it is evident that $\hat{S}_{HSR,EP}$ will be least sensitive to error in $\hat{\mu}$ when HSR = 1. This condition is satisfied along the ridges in the right panels of the figure. Error in $\hat{\gamma}$ dominates error in $\hat{S}_{HSR,EP}$ when μ is large; hence $\hat{S}_{HSR,VSR}$ is more reliable than $\hat{S}_{HSR,EP}$ at near distances.

7.2 *Slant-estimation vs slant-nulling tasks*

We are now in a position to explain why the induced effect increases more with distance in the nulling than in the estimation task. With a vertical magnifier before one eye, $\hat{S}_{HSR,VSR}$ and $\hat{S}_{HSR,EP}$ disagree. In the estimation task, $\hat{S}_{HSR,VSR}$ is nonzero and scales with distance, but $\hat{S}_{HSR,EP}$ is always zero. $\hat{S}_{HSR,VSR}$ decreases in reliability more rapidly than does $\hat{S}_{HSR,EP}$, so the visual system gives relatively less weight to $\hat{S}_{HSR,VSR}$ as distance increases. As a result, $\hat{S}_{stereo}$ does not scale with distance. In the slant-nulling task, $\hat{S}_{HSR,EP}$ and $\hat{S}_{HSR,VSR}$ are not fixed, but rather change as the surface is adjusted to appear gaze normal. It is when $\hat{S}_{HSR,VSR}$ equals 0° (along the ridge in the lower left panel of figure 11) that $\hat{S}_{HSR,VSR}$ is most reliable. Thus $\hat{S}_{stereo}$ will resemble $\hat{S}_{HSR,VSR}$ more in the nulling than in the estimation task, and as a result, slant settings increase with distance more in the nulling task.

7.3 *Nonstereo cues*

At this point, it is necessary to reconsider the role of nonstereo cues because they offer a potentially useful alternative explanation. $S_{nonstereo}$ was 0° in our experiments and simulations. Because the final perceived slant was 0° in the nulling task, $S_{nonstereo}$ cannot affect $\hat{S}$ in that task. However, $S_{nonstereo}$ can affect $\hat{S}$ in the estimation task. The reliability of stereoscopic cues decreases with distance while the reliability of nonstereo cues does not (Sedgwick 1986). As a consequence, increasing reliance on $S_{nonstereo}$ as distance increases would explain why the induced effect scaled with distance for nulling but not for estimation.

There are three reasons why the use of nonstereo cues cannot account for all of the observed difference between the estimation and nulling tasks. First, a real surface was used in the nulling experiments of Ogle (1938), so $S_{nonstereo}$ was equal to $S_{HSR,EP}$ at the final slant settings; still, the induced effect scaled with distance. Second, both the model and experiments confirmed that the difference between estimation and nulling is less dramatic for horizontal than for vertical magnification. Increased reliance on $S_{nonstereo}$ cannot by itself explain this difference between the geometric and induced effects because it is put in conflict with stereoscopic signals in both. Finally, the results of experiment 3 showed that the conflict between sensed eye position and vertical disparities does indeed contribute to the failure of distance scaling in the induced effect for slant-estimation tasks. Thus, nonstereo cues cannot account for all of the difference between tasks. Nevertheless, nonstereo cues account for much of this difference; indeed, inspection of figure 13 shows how, in our model, all of the intertask difference for horizontal magnification results from the decreasing reliability of stereo cues as distance increases.

7.4 *The data of Amigo*

Amigo (1972) conducted a slant-nulling experiment and reported a small induced effect that actually decreased with distance. This result is inconsistent with Ogle (1938) and our experiment 2. However, Amigo's stimuli consisted of a horizontal row of small beads, flanked above and below by single small beads. It is surely more difficult for the visual system to measure VSR in this stimulus than in large random-dot stimuli (Ogle 1939; Backus et al 1999). This would reduce reliance on $S_{HSR,VSR}$ relative to $S_{HSR,EP}$ and the latter two indicated zero slant when objective slant was zero. When our simulation is rerun with a larger error in $\widehat{VSR}$, the model also reports a decrease in the objective slant of null settings as distance increases.

7.5 *Overestimation of slant responses by the model*

The Monte Carlo simulation overestimated slant responses in the estimation tasks (experiments 1 and 3). There are several possible reasons for the overestimation. First, we may have underestimated the reliability of perspective and other nonstereo cues that indicated zero slant. Second, observers may have a perceptual bias toward zero slant. Finally, we assumed that slant responses in the estimation task are equal to internal slant estimates. This assumption may have been incorrect; perhaps responses were simply smaller than internal representations. In general, it is easiest to predict data for a procedure that can be modeled without knowing how internal states are related to responses. Banks and Dannemiller (1987) distinguished between this type of procedure, called an *input-mapping* procedure, and a procedure for which a model must include the mapping of internal states onto responses, called an *output-mapping* procedure. Slant nulling is an input-mapping procedure and slant estimation is an output-mapping procedure.

7.6 *Why are individual differences so prominent in these tasks?*

There were large differences between observers' slant responses, particularly in experiments 1 and 3. Why is the interobserver variability so large in induced-effect and geometric-effect experiments? Part of the answer must be that the stimuli created in these experiments can produce striking cue conflicts because one does not observe the same degree of interobserver variability when cue conflicts are reduced: for example, in the experiments of Backus et al (1999), we found that interobserver variability was smaller when $\hat{S}_{\mathrm{HSR,VSR}}$ and $\hat{S}_{\mathrm{HSR,EP}}$ agreed (experiment 1) than when they did not agree (experiment 5). Thus, the question becomes why variability between observers is larger in cue-conflict situations.

From the perspective of our model (figure 10), there are a number of ways that observers can differ: (1) noises of the signal estimates, (2) biases in the signal estimates, (3) weights given to the slant estimator outputs, and (4) mapping from the final slant estimate to the behavioral response. In our simulations, we assumed that there were no biases in signal estimates (except in gain_{μ}, figure 16), that the weights given to estimator outputs were optimal, and that the mapping from final slant estimate to response was veridical. The fourth potential difference can plausibly account for order-preserving differences in responses, but we observed some non-order-preserving changes; for example, there were cases in which slant responses increased with distance for some observers and decreased for others. Moreover, changes in the mapping between estimate and response would not explain the larger individual differences in cue-conflict as compared with natural-viewing situations (eg experiment 1 vs experiment 5 in Backus et al 1999). Thus, difference in the estimate-to-response mapping is by itself an unlikely explanation for the large interobserver differences observed in cue-conflict situations like the induced and geometric effects.

7.7 *Are estimator reliabilities learned or calculated?*

The reliability of each slant estimator, as determined by the visual system, should be a function of the viewing situation. For instance, the reliability of $\hat{S}_{\mathrm{HSR,VSR}}$ ought to be higher for short viewing distances than for long. There are two interesting issues concerning how estimator reliabilities are established.

First, are the reliabilities optimal as they would be for an ideal observer? Our simulations employed optimal reliabilities, given our assumptions, and the fact that the simulations yielded behavior rather similar to human behavior in the same viewing situations suggests that the estimated reliabilities may approach optimality.

Second, how are the reliabilities of slant estimators acquired? Two general acquisition methods could work: (i) they could be learned over time by comparing the outputs of individual estimators with those of other estimators and even with different systems

7.5 *Overestimation of slant responses by the model*

The Monte Carlo simulation overestimated slant responses in the estimation tasks (experiments 1 and 3). There are several possible reasons for the overestimation. First, we may have underestimated the reliability of perspective and other nonstereo cues that indicated zero slant. Second, observers may have a perceptual bias toward zero slant. Finally, we assumed that slant responses in the estimation task are equal to internal slant estimates. This assumption may have been incorrect; perhaps responses were simply smaller than internal representations. In general, it is easiest to predict data for a procedure that can be modeled without knowing how internal states are related to responses. Banks and Dannemiller (1987) distinguished between this type of procedure, called an *input-mapping* procedure, and a procedure for which a model must include the mapping of internal states onto responses, called an *output-mapping* procedure. Slant nulling is an input-mapping procedure and slant estimation is an output-mapping procedure.

7.6 *Why are individual differences so prominent in these tasks?*

There were large differences between observers' slant responses, particularly in experiments 1 and 3. Why is the interobserver variability so large in induced-effect and geometric-effect experiments? Part of the answer must be that the stimuli created in these experiments can produce striking cue conflicts because one does not observe the same degree of interobserver variability when cue conflicts are reduced: for example, in the experiments of Backus et al (1999), we found that interobserver variability was smaller when $\hat{S}_{\mathrm{HSR,VSR}}$ and $\hat{S}_{\mathrm{HSR,EP}}$ agreed (experiment 1) than when they did not agree (experiment 5). Thus, the question becomes why variability between observers is larger in cue-conflict situations.

From the perspective of our model (figure 10), there are a number of ways that observers can differ: (1) noises of the signal estimates, (2) biases in the signal estimates, (3) weights given to the slant estimator outputs, and (4) mapping from the final slant estimate to the behavioral response. In our simulations, we assumed that there were no biases in signal estimates (except in gain_{μ}, figure 16), that the weights given to estimator outputs were optimal, and that the mapping from final slant estimate to response was veridical. The fourth potential difference can plausibly account for order-preserving differences in responses, but we observed some non-order-preserving changes; for example, there were cases in which slant responses increased with distance for some observers and decreased for others. Moreover, changes in the mapping between estimate and response would not explain the larger individual differences in cue-conflict as compared with natural-viewing situations (eg experiment 1 vs experiment 5 in Backus et al 1999). Thus, difference in the estimate-to-response mapping is by itself an unlikely explanation for the large interobserver differences observed in cue-conflict situations like the induced and geometric effects.

7.7 *Are estimator reliabilities learned or calculated?*

The reliability of each slant estimator, as determined by the visual system, should be a function of the viewing situation. For instance, the reliability of $\hat{S}_{\mathrm{HSR,VSR}}$ ought to be higher for short viewing distances than for long. There are two interesting issues concerning how estimator reliabilities are established.

First, are the reliabilities optimal as they would be for an ideal observer? Our simulations employed optimal reliabilities, given our assumptions, and the fact that the simulations yielded behavior rather similar to human behavior in the same viewing situations suggests that the estimated reliabilities may approach optimality.

Second, how are the reliabilities of slant estimators acquired? Two general acquisition methods could work: (i) they could be learned over time by comparing the outputs of individual estimators with those of other estimators and even with different systems

such as visually guided reaching and grasping, and (ii) they could be estimated directly from some measure of the estimator output itself such as its fluctuation over time. It is important to note, however, that in either case the nervous system has to make assumptions for the acquisition method to proceed. In the case of learning over time, it would have to assume which estimator outputs should agree and how they should interface with visuomotor behavior. In the case of estimating reliability directly, the nervous system would have to assume that the environmental signals that affect the slant estimator are not changing (eg that the surface is not rotating and thereby affecting HSR).

In the context of learning over time, Landy et al (1995) define an 'ancillary measure' as a signal by which the visual system recognizes the viewing situation, and hence the reliability of a cue. The analogy is with 'ancillary statistics', which are independent of the value of the parameter being estimated (Landy et al 1995). We see no reason why the visual system could not use the same signals both to compute slant and to estimate the reliability of the resulting computation. We suggest that a distinction be made between methods of measuring reliability in general, and ancillary measures in particular.

7.8 *Conclusion*

When one eye's image is magnified vertically, slant-nulling tasks (experiment 2; Ogle 1938) reveal a clear effect of distance on perceived slant, but slant-estimation tasks exhibit little effect of distance (experiment 1; Gillam et al 1988; Rogers et al 1995). An analysis of slant estimator reliability in the presence of noise resolves the apparent contradiction. Our analysis depends on a quantitative description of the relationship between environmental signals and surface slant that can be described by three slant estimators—$\hat{S}_{\mathrm{HSR,\,VSR}}$, $\hat{S}_{\mathrm{HSR,\,EP}}$, and $\hat{S}_{\mathrm{nonstereo}}$—and the assumption that the visual system weights the outputs of these estimators according to their reliability even when the estimates greatly conflict. The analysis also predicts that if eye position is manipulated, the distance scaling normally observed in forward-gaze experiments can be reversed such that apparent slant should increase more with distance under vertical than under horizontal magnification. This prediction was confirmed experimentally.

Acknowledgements. This work was supported by research grants from AFOSR (93NL366) and NSF (DBS-9309820) and by the Max Planck Institute for Biological Cybernetics, Tübingen, Germany. We thank Sarah Freeman, Terri Yamamoto, Jamie Hillis, and Winnie Fan for participating as observers.

References

Amigo G, 1972 "The stereoscopic frame of reference in asymmetric convergence of the eyes: Response to 'point' stimulation of the retina" *Optica Acta* **19** 993 – 1006

Backus B T, Banks M S, Ee R van, Crowell J A, 1999 "Horizontal and vertical disparity, eye position, and stereoscopic slant perception" *Vision Research* **39** in press

Banks M S, Backus B T, 1998 "Extra-retinal and perspective cues cause the small range of the induced effect" *Vision Research* **38** 187 – 194

Banks M S, Dannemiller J L, 1987 "Visual psychophysics", in *Handbook of Infant Perception* Eds P Salapatek, L Cohen (New York: Academic Press) pp 115 – 184

Blake A, Bülthoff H H, Sheinberg D, 1993 "Shape from texture: ideal observers and human psychophysics" *Vision Research* **33** 1723 – 1737

Bradshaw M F, Glennerster A, Rogers B J, 1996 "The effect of display size on disparity scaling from differential perspective and vergence cues" *Vision Research* **36** 1255 – 1264

Brenner E, Damme W J M van, 1998 "Judging distance from ocular convergence" *Vision Research* **38** 493 – 498

Buckley D, Frisby J P, 1993 "Interaction of stereo, texture and outline cues in the shape perception of three-dimensional ridges" *Vision Research* **33** 919 – 933

Buell T N, Hafter E R, 1991 "Combination of binaural information across frequency bands" *Journal of the Acoustical Society of America* **90** 1894 – 1900

Bülthoff H H, Mallot H A, 1988 "Integration of depth modules: stereo and shading" *Journal of the Optical Society of America A* **5** 1749 – 1758

Clark J J, Yuille A L, 1990 *Data Fusion for Sensory Information Processing Systems* (Boston, MA: Kluwer Academic)

Dosher B A, Sperling G, Wurst S A, 1986 "Tradeoffs between stereopsis and proximity luminance covariance as determinants of perceived 3D structure" *Vision Research* **26** 973 – 990

Ee R van, Erkelens C J, 1996 "Temporal aspects of binocular slant perception" *Vision Research* **36** 43 – 51

Foley J M, 1980 "Binocular distance perception" *Psychological Review* **87** 411 – 434

Gårding J, Porrill J, Mayhew J E, Frisby J P, 1995 "Stereopsis, vertical disparity and relief transformations" *Vision Research* **35** 703 – 722

Gillam B, Chambers D, Lawergren B, 1988 "The role of vertical disparity in the scaling of stereoscopic depth perception: An empirical and theoretical study" *Perception & Psychophysics* **44** 473 – 483

Gillam B, Lawergren B, 1983 "The induced effect, vertical disparity, and stereoscopic theory" *Perception & Psychophysics* **34** 121 – 130

Heller L M, Trahiotis C, 1996 "Extents of laterality and binaural interference effects" *Journal of the Acoustical Society of America* **99** 3632 – 3637

Huber P J, 1981 *Robust Statistics* (New York: John Wiley)

Johnston E B, Cumming B G, Parker A J, 1993 "Integration of depth modules: Stereopsis and texture" *Vision Research* **33** 813 – 826

Kaneko H, Howard I P, 1997 "Spatial limitation of vertical-size disparity processing" *Vision Research* **37** 2871 – 2878

Landy M S, Maloney L T, Johnston E B, Young M, 1995 "Measurement and modeling of depth cue combination: In defense of weak fusion" *Vision Research* **35** 389 – 412

Mayhew J E W, Longuet-Higgins H C, 1982 "A computational model of binocular depth perception" *Nature (London)* **297** 376 – 378

Ogle K N, 1938 "Induced size effect. I. A new phenomenon in binocular space-perception associated with the relative sizes of the images of the two eyes" *Archives of Ophthalmology* **20** 604 – 623

Ogle K N, 1939 "Induced size effect. II. An experimental study of the phenomenon with restricted fusion stimuli" *Archives of Ophthalmology* **21** 604 – 625

Rogers B J, Bradshaw M F, 1993 "Vertical disparities, differential perspective and binocular stereopsis" *Nature (London)* **361** 253 – 255

Rogers B J, Bradshaw M F, 1995 "Disparity scaling and the perception of frontoparallel surfaces" *Perception* **24** 155 – 179

Rogers B J, Bradshaw M F, Gillam B, 1995 "The induced effect does not scale with viewing distance" *Perception* **24** Supplement, 33

Schor C M, Maxwell J S, Stevenson S B, 1994 "Isovergence surfaces: the conjugacy of vertical eye movements in tertiary positions of gaze" *Ophthalmology and Physiological Optics* **14** 279 – 286

Sedgwick H, 1986 "Space perception", in *Handbook of Perception and Human Performance* volume I *Sensory Processes and Perception* Eds K R Boff, L Kaufman, J P Thomas (New York: John Wiley) pp 21-1 – 21-57

Stevens K A, Brookes A, 1988 "Integrating stereopsis with monocular interpretations of planar surfaces" *Vision Research* **28** 371 – 386

Wallach H, Moore M E, Davidson L, 1963 "Modification of stereoscopic depth perception" *American Journal of Psychology* **76** 191 – 204

APPENDIX

Details of the simulations

We assumed that noise in signal measurement was Gaussian with mean 0. The standard deviations we assumed for the measurements of ln HSR, ln VSR, μ, γ, and nonstereo slant ($\widehat{\ln \text{HSR}}$, $\widehat{\ln \text{VSR}}$, $\hat{\mu}$, $\hat{\gamma}$, and $\hat{S}_{\text{nonstereo}}$) were respectively, $\varepsilon_{\ln \text{HSR}} = 1.3 \times 10^{-3}$, $\varepsilon_{\ln \text{VSR}} = 1.3 \times 10^{-3}$, $\varepsilon_{\mu} = 0.5°$, $\varepsilon_{\gamma} = 1°$, and $\varepsilon_{\hat{S}_{\text{nonstereo}}} = 5°$.[9,10] The results of the simulation are qualitatively similar if any of these noise values are multiplied or divided by 2. The assumption that $\hat{\mu}$ is normally distributed in the visual system is probably false when μ is small: a sensible mechanism would not accept $\hat{\mu} < 0$. In the simulation, we replaced samples of $\hat{\mu}$ that were smaller than 0.5° with the value 0.5°. At distance = 172 cm, $\mu = 2.0°$, so the minimum allowed value was 3 standard deviations from μ. Thus, 0.1% of trials were affected at that distance. No samples were replaced at other distances.

For each viewing condition considered, the simulation drew a value from each signal measurement distribution, and calculated the three slant estimates. From 40 000 trials, we determined the distributions of the two stereoscopic slant estimators and the nonstereo slant estimator. The two stereoscopic estimators are not independent, so on each trial their outputs were combined (before combination with $\hat{S}_{\text{nonstereo}}$) by using weights proportional to $r_{\hat{S}_{\text{HSR, VSR}}}$ and $r_{\hat{S}_{\text{HSR, EP}}}$ (calculated a posteriori from the distributions). The reliability of the resulting stereoscopic slant estimator $\hat{S}_{\text{stereo}}$ is calculated from the variance (over trials) of $\hat{S}_{\text{stereo}}$.[11] The final slant estimate $\hat{S}$ is then calculated as the weighted average of $\hat{S}_{\text{stereo}}$ and $\hat{S}_{\text{nonstereo}}$. This weighted average was the model's output.

[9] At this point we know little about the precision with which the signals considered here can be measured by the visual system. Thus, the assumed standard deviations are somewhat arbitrary. The values for $\varepsilon_{\ln \text{HSR}}$ and $\varepsilon_{\ln \text{VSR}}$ came from the observation that a change of 0.2% in HSR is just visible as a slant change in the nulling task; this provides an upper bound on total signal noise in equation (1) for 0° slant. We also assumed that $\varepsilon_{\ln \text{HSR}} = \varepsilon_{\ln \text{VSR}}$. The value we chose for $\varepsilon_{\ln \text{VSR}}$ may be an overestimate because VSR can be measured by averaging over a large region of the visual field and HSR appears to be measured locally (Kaneko and Howard 1997). The estimate of ε_{γ} comes from an unpublished pointing experiment from our laboratory. The estimate of ε_{μ} derives from the assumption that $\varepsilon_{\mu} < \varepsilon_{\gamma}$: although a difference in measured eye position (μ) is in principle noisier than an average (γ), the range of μ is smaller than that of γ, and μ can be estimated from accommodation and knowledge of our apparatus, which γ cannot. The assumed value of ε_{μ} is also consistent with recent observations by Brenner and van Damme (1998). $\varepsilon_{\hat{S}_{\text{nonstereo}}}$ is based on results from the experiment described in section 6.1. With a real surface, the reliability of $\hat{S}_{\text{nonstereo}}$ might vary with $S_{\text{objective}}$. We assumed constant reliability for $\hat{S}_{\text{nonstereo}}$ because in our experiments the perspective slant cue did not vary with objective slant.

[10] As noted earlier, μ can be measured either from eye position or the retinal images. In forward gaze, μ is equal to $\partial \text{VSR}/\partial\gamma$ in the images. As with the other signals, we assume that noise in the measurement of this signal is independent of its value. In eccentric gaze, μ also depends on HSR and VSR. This could add a small amount of additional noise to the retinal measurement of μ in eccentric gaze [see equation (4) in Backus et al (1999)].

[11] Variance in the average of two random variables is independent of their means. Thus, the reliability of $\hat{S}_{\text{stereo}}$ does not automatically go down for cue-conflict stimuli as might be supposed.