An analysis of binocular slant contrast

Raymond van Ee[1,2], Martin S. Banks[1,3], Benjamin T. Backus[1,4]

[1] *School of Optometry and Vision Science Program, University of California, Berkeley*

[2] *Helmholtz Institute, Princetonplein 5. 3584 CC, Utrecht, The Netherlands*

[3] *Department of Psychology, University of California, Berkeley*

[4] *Department of Psychology, Stanford University*

Abstract. When a small frontoparallel surface (a test strip) is surrounded by a larger slanted surface (an inducer), the test strip is perceived as slanted in the direction opposite to the inducer. This has been called the depth-contrast effect, but we call it the slant-contrast effect. In nearly all demonstrations of this effect, the inducer's slant is specified by stereoscopic signals, and other signals, such as the texture gradient, specify that it is frontoparallel. We present a theory of slant estimation that determines surface slant via linear combination of various slant estimators; the weight of each estimator is proportional to its reliability. The theory explains slant contrast because the absolute slant of the inducer and the relative slant between test strip and inducer are both estimated with greater reliability than the absolute slant of the test strip. The theory predicts that slant contrast will be eliminated if the signals specifying the inducer's slant are consistent with one another. It also predicts reversed slant contrast if the inducer's slant is specified by nonstereoscopic signals rather than by stereo signals. These predictions were tested and confirmed in three experiments. The first showed that slant contrast is greatly reduced when the stereo- and nonstereo-specified slants of the inducer are made consistent with one another. The second showed that slant contrast is eliminated altogether when the stimulus consists of real planes rather than images on a display screen. The third showed that slant contrast is reversed when the nonstereo-specified slant of the inducer varies and the stereo-specified slant is zero. We conclude that slant contrast is a byproduct of the visual system's reconciliation of conflicting information while it attempts to determine surface slant.

1 Introduction

There is a variety of stereoscopic phenomena in which the perceived slant or curvature of a surface patch is not what one would predict from the disparities. A well-known example is the depth-contrast effect which is illustrated in Figure 1. The half image of the cross-hatched plane on the right is wider than the one on the left. When fused, the planes create the disparity gradient associated with a plane rotated about a vertical axis. The half images of the strip of dots are identical, so the strip creates the disparity gradient associated with a frontal plane. When fused, however, the perceived slant of the strip of dots is not frontoparallel; rather it appears rotated in the direction opposite from the specified slant of the cross-hatched planes. The illusory slant is the depth-contrast effect. It is actually an illusion of slant[1], so we will refer to it as the *slant-contrast* effect[2] in the remainder of the paper. Our purpose here is to explain slant contrast in the context of current theory.

Werner (1937, 1938) investigated the slant-contrast effect six decades ago. He pointed out that the perceived slant of the surrounding stimulus (which we refer to as the *inducer*) was less than the slant specified by the disparity gradient and that the perceived slant of the internal stimulus (the *test strip*) was opposite from the disparity-specified slant of the inducer. In the past 60 years, slant contrast has been reported for a very wide variety of stimuli (Ogle, 1946; Harker, 1962; Pastore, 1964; Pastore & Terwilliger, 1966; Nelson, 1977; Graham & Rogers, 1982; Stevens & Brookes, 1988; Fahle & Westheimer, 1988; Brookes & Stevens, 1989; Howard, Ohmi, & Sun, 1993; van Ee & Erkelens, 1996b; Pierce, Howard, & Feresin, 1998; Sato & Howard, 1999).

---

[1] To describe the orientation of a surface patch, we use Stevens' (1983) convention. Slant is the angle between the surface normal and the line of sight. Tilt denotes the slant direction which is the angle between the projection of the surface normal onto the frontoparallel plane and the horizontal in the frontal plane. Here we examine slant about the vertical axis (tilt = 0˚) and about the horizontal axis (tilt = 90˚). We consider forward gaze only. Slant is defined relative to the head, so zero slant means parallel to the forehead.

[2] We have chosen the phrase *slant contrast* in order to make a distinction between the phenomenon examined here and the phenomenon of *depth contrast* which concerns misperceptions of distance caused by the interaction between one target and another (e.g., Richards, 1972; Anstis, 1975; Mitchison & Westheimer, 1984; Westheimer & Levi, 1987; Stevenson, Cormack, & Schor, 1991; Kumar & Glaser, 1991, 1993).

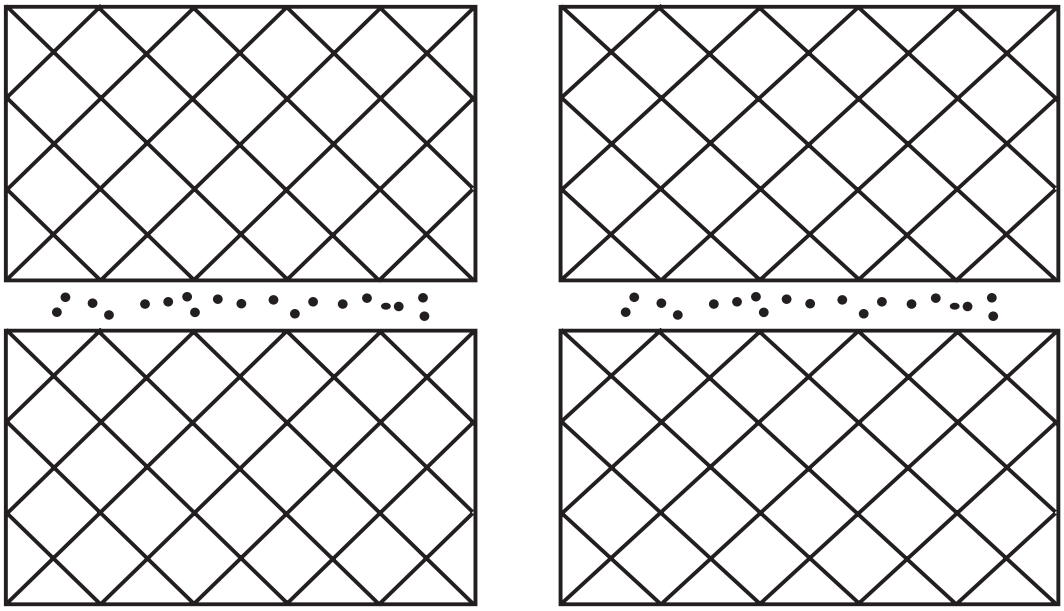

Figure 1. The slant-contrast effect. The crossed-hatched patterns are the *inducer* and the strip of random dots is the *test strip*. When you cross fuse this image, the test strip is identical in the two eyes, so it should appear frontoparallel. The inducer is 6% wider in the left eye (the opposite for divergent fusing), so it should appear rotated about a vertical axis such that its left side is more distant than its right side (opposite for divergent fusing). For most people, the test strip appears slanted such that its left side is nearer than its right side; this is the slant-contrast effect.

There are many explanations for slant contrast (see Howard & Rogers, 1995, for a recent review), but they all should explain two aspects of the phenomenon: underestimation[3] of the inducer's slant and the induced slant of the test strip.

A variety of theories attempt to explain underestimation of the inducer's slant. They involve a re-mapping of the function that relates disparity gradients to perceived slant such that the inducer is perceived as less slanted than predicted from the disparity gradient.

One theory states that normalization[4] is based on disparity alone. It proposes that the visual system is less sensitive to the disparity gradient associated with a slanted plane (so-called first-order disparity) than it is to the disparity gradient between two planes (higher-order disparity; Gillam, Flagg, & Finlay, 1984; Stevens & Brookes, 1988; van Ee & Erkelens, 1996c). The induced slant of the test strip is due to preservation or enhancement of its slant *relative* to the inducer (van Ee & Erkelens, 1996b). Center-surround, antagonistic mechanisms have been proposed that would have the behavior of preserving or enhancing slant relative to the inducer (Anstis, Howard & Rogers, 1978; Schumer & Ganz, 1979; Rogers & Graham, 1983; Mitchison & Westheimer, 1984; Brookes & Stevens, 1989).

Another theory, differing slightly from the first, emphasizes normalization caused by a perceptual bias to see surfaces as frontal (Gogel, 1956; Harker, 1962; Howard & Rogers, 1995). According to this view, any surface would be perceived as less slanted than specified by the disparity gradient. Although

[3] Underestimation is an unfortunate term for two reasons. First, observers presumably estimate slant to the best of their ability given the available information. Indeed, with the stimuli used in slant-contrast experiments, the inducer's slant is underestimated with respect to its disparity-specified slant and not with respect to the slant specified by other cues such as the texture gradient. Second, in these experiments observers indicate the perceived slant by setting a line's orientation in a symbolic display (e.g., van Ee & Erkelens, 1996a), by setting a handplate to the perceived orientation (e.g., Howard & Kaneko, 1994), or by some other method. In all cases, the settings are determined not only by the perceived slant but also by the function that maps percepts onto responses. Because we do not know the form of that mapping function, there are no grounds by which to determine what set of responses would indicate veridical percepts.

[4] In normalization, a stimulus that is to one side of a norm on a sensory dimension appears more similar to the norm than it is. Slanted surfaces, therefore, ought to appear less slanted than they are because the norm is the frontal plane.

this theory might explain underestimation of the inducer's slant, it is unclear how it can explain slant contrast. In particular, it is unclear how a bias toward the frontal plane can by itself explain the fact that the test strip would appear slanted when it is actually frontoparallel. We can, therefore, reject this theory as an adequate account for the slant-contrast phenomenon.

These first two sets of theories are inconsistent with the observation that the visual system is actually quite sensitive to first-order disparity per se: When nonstereoscopic slant cues are uninformative, observers can adjust the slant of a random-dot plane with a standard deviation in disparity gradient of about 0.2%, or ~7 arcsec per degree of visual angle (Ogle, 1938; Backus, Banks, van Ee, & Crowell, 1999; Backus & Banks, 1999). This precision compares favorably with the disparity threshold of 3–10 arcsec for a depth edge (Howard & Rogers, 1995). We will discuss why in previous experiments the visual system appears to be relatively insensitive to first-order disparity when we present our model.

Another slant-contrast theory is based on nonstereoscopic cues. It claims that conflicting nonstereo cues, such as the outline shape of the inducer, cause a remapping of the function between disparity gradients and perceived slant (Kumar & Glaser, 1991, 1992, 1993) such that a nonzero disparity gradient comes to be interpreted as frontoparallel by the observer. We will show that remapping is not necessary to explain underestimation and slant contrast.

Ogle (1946) claimed, without providing data, that slant contrast did not occur if the inducer had the outline shape of a real slanted rectangle. This claim suggests that slant contrast is eliminated when the inducer's disparity-specified and perspective-specified slants are consistent with one another.[5] Everyday experience is also consistent with Ogle's claim because stimuli that produce slant contrast are difficult to find in the natural environment. Theories based on relative insensitivity to first-order disparities (Anstis et al, 1978; Rogers & Graham, 1983; van Ee & Erkelens, 1996b) did not incorporate nonstereoscopic cues in their explanations, so they would predict no effect of manipulating the perspective-specified slant of the inducer.

2 Slant Estimation Theory

The standard slant-contrast stimulus presents a variety of signals to the visual system, so to analyze the slant-contrast effect, it should be profitable to describe those signals and then consider their perceptual consequences. We first define some terms and then turn to the analysis of the effect. A *signal* is a real-valued parameter that characterizes an aspect of the viewing situation and might be measured by the nervous system. A *signal measurement* is the representation of a signal within the nervous system; the measurement inevitably contains some error. An *estimator* is a mechanism that combines one or more measurements to yield an *estimate* of a scene parameter; different estimators may provide estimates of the same scene parameter using different signal measurements. For the present purposes, we will consider slant estimators only. The *reliability* of an estimator is the reciprocal variance of its output for a given viewing situation (Backus & Banks, 1999). We assume that the estimator reliabilities are known approximately by the visual system, and that the system uses those reliabilities in a statistically sensible way while combining the estimator outputs (Clarke & Yuille, 1990; Landy, Maloney, Johnston, & Young, 1995; Backus & Banks, 1999).

A surface cannot have two slants at the same time, so the visual system must construct a unique slant estimate for each surface patch in the scene. We assume that this construction involves reconciliation of many estimator outputs that can differ from one another. As we will show, the standard slant-contrast stimulus presents inconsistent signals and thus the analysis of the slant-contrast effect hinges on the manner by which the inconsistency is reconciled.

Consider the two parts of the stimulus in isolation. The physical slants of the inducer and test strip are S_i and S_t. We assume that the perceived slants of the isolated inducer and isolated test strip are given by linear combination of weighted contributions from stereo and nonstereo estimators (Backus & Banks, 1999; Landy et al, 1995).[6]

---

[5] Kumar and Glaser (1992) showed that the inducer's outline shape influenced the perceived slant of a test stimulus and suggested that slant contrast was dependent on interactions between stereoscopic and perspective information; they did not, however, explicate the nature of those interactions.

[6] Notice in equation (1) and subsequent equations that a change in the weight given to an individual slant estimator (e.g., the weight w_{id} given to the estimator $\hat{S}_{id}$) can be exactly reproduced by a change in the estimator's gain. It is beyond the purposes of this paper to examine the distinction between a weight change and an estimator gain change, but we will take up this issue in a forthcoming paper.

$$\hat{S}_{i,alone} = w_{id}\hat{S}_{id} + w_{ip}\hat{S}_{ip} \quad (1A)$$

$$\hat{S}_{t,alone} = w_{td}\hat{S}_{td} + w_{tp}\hat{S}_{tp} \quad (1B)$$

$\hat{S}$ is the output from a slant estimator and w is a weight; the subscripts i and t refer to the inducer and test strip, respectively, and the subscripts d and p to stereo and nonstereo (d for disparity and p for perspective), respectively. The weights are positive and add to 1 (e.g., $w_{id} + w_{ip} = 1$). The stereo estimates, $\hat{S}_{td}$ and $\hat{S}_{id}$, are based on the gradients of horizontal disparity created by the test strip and inducer (plus other signals that correct for distance and azimuth; Backus et al, 1999). Similarly, the nonstereo estimates, $\hat{S}_{tp}$ and $\hat{S}_{ip}$, are based on perspective cues created by the strip and inducer (e.g., foreshortening of texture elements and projected outline shape; Sedgwick, 1986).

We need also to consider signals created jointly by the two parts of the stimulus. S_{rel} ("relative slant") is the difference between the slants of the two surfaces, as specified by disparity. $\hat{S}_{rel}$ is based on the disparity gradient between the two surfaces (plus other signals needed to correct for distance and azimuth). Such relative disparity gradients are very effective in signaling that adjacent surfaces have different slants (Gillam et al, 1984; van Ee & Erkelens, 1996c).[7] If all three stereo estimators ($\hat{S}_{id}$, $\hat{S}_{td}$, and $\hat{S}_{rel}$) are unbiased, it will on average be the case that:

$$\hat{S}_{td} = \hat{S}_{id} + \hat{S}_{rel} \quad (2)$$

The process of estimating the inducer's and test strip's slants from the signals they create is depicted in the upper panels of Figure 2. The left panel portrays the estimation of the test strip's slant when presented in isolation; the stereo- and nonstereo-specified slants, S_{td} and $\hat{S}_{tp}$, are both zero, so the slant estimators, $\hat{S}_{td}$ and $\hat{S}_{tp}$, on average specify slants of zero. Therefore, the slant estimate, $\hat{S}_{t,alone}$, is zero. The upper middle panel portrays the estimation of the inducer's slant when presented in isolation. The stereo-specified slant, S_{id}, is represented by the black dashed line and the nonstereo slant, S_{ip}, by the gray dashed line. The two estimates differ because the nonstereo-specified slant is always zero and the stereo slant is non-zero. The solid line represents the final slant estimate, $\hat{S}_{i,alone}$, which is a compromise between the stereo- and nonstereo-specified slants. The right panel represents the stereoscopic measurement $\hat{S}_{rel}$.

In the natural environment, signals created by surfaces in isolation and by the signal created between surfaces are on average consistent with one another. In the slant-contrast stimulus, however, they are inconsistent. To see how, first consider the signals from the surfaces in isolation (equation (1A,B)): the estimate for the inducer in isolation ($\hat{S}_{i,alone}$) will be different from both its stereo-based and nonstereo-based estimates: $\hat{S}_{i,alone} \neq \hat{S}_{id} \neq \hat{S}_{ip}$. In the upper middle panel of Figure 2, $\hat{S}_{id} = -30°$, $\hat{S}_{ip} = 0°$, so $\hat{S}_{i,alone} \approx -15°$ (assuming equal weights for the stereo- and nonstereo-specified slants). The estimated slant difference derived from the isolated surfaces is $\hat{S}_{t,alone} - \hat{S}_{i,alone}$, which would be 15° in the figure. Now consider the relative slant specified by the disparity gradient between the two surfaces: the estimated slant of one surface should differ from the estimated slant of the other by $\hat{S}_{rel}$. In Figure 2 the difference would be 30°. Thus, the slant difference specified by the isolated signals is not the same as that specified by the joint signal. Somehow the visual system must reconcile the disagreement. We assume that it does so by combining information from the available signals and weighting each information source according to its estimated reliability. For the test strip:

$$\hat{S}_t = w_{t,dir}\hat{S}_{t,alone} + w_{t,ind}(\hat{S}_{i,alone} + \hat{S}_{rel}) \quad (3A)$$

[7] One could include a nonstereo component to this estimator, but we do not because even if direct measurement of the difference in perspective gradients does occur, it would not be very reliable for the stimuli in question because the test strip's perspective cues are weak.

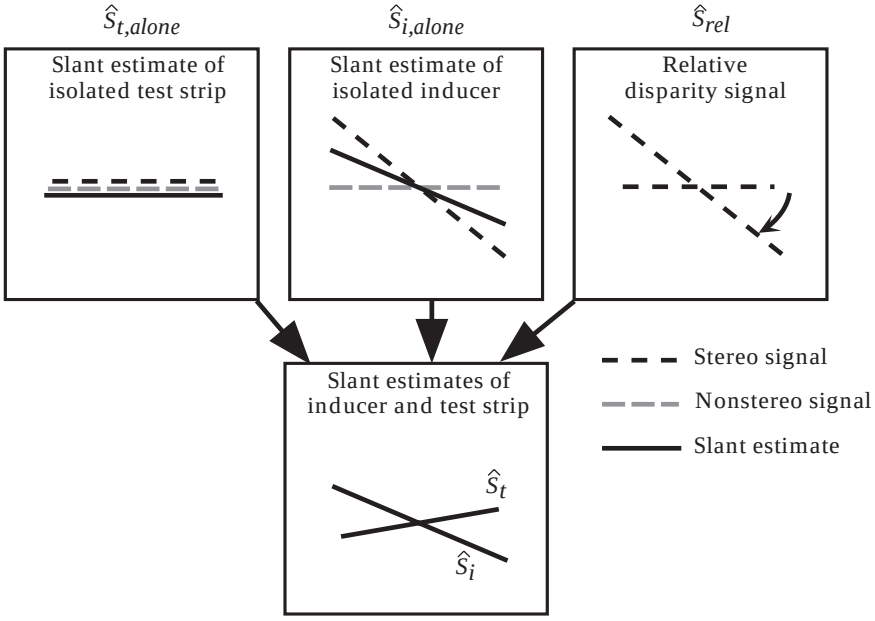

Figure 2. Slant estimation when the inducer's nonstereo-specified slant is inconsistent with its stereo-specified slant. Most demonstrations of slant contrast present signals like those depicted here. The panels represent the slant signals arising from the test strip, inducer, and the relationship between them; they also represent the means by which the slants of the test strip and inducer are estimated according to the slant estimation theory described in the text. Upper left panel: slant estimation for the isolated test strip. The black dashed line represents the stereo-specified slant, the gray dashed line the nonstereo-specified slant, and the solid black line the slant estimate once the stereo and nonstereo estimates are combined (equation 1B). Upper middle panel: slant estimation for the isolated inducer. The black and gray dashed lines represent the stereo- and nonstereo-specified slants, respectively, and the black solid line represents the slant estimate once the stereo and nonstereo estimates are combined in a weighted average (equation 1A). Upper right panel: estimation of the slant difference between the inducer and test strip from the relative disparity gradient between the two surfaces. Bottom panel: final slant estimates for the test strip and inducer. The solid black lines represent the slant estimates for the two parts of the stimulus assuming that the inducer's stereo- and nonstereo-specified slants are given roughly equal weight. The test strip's slant is estimated in the fashion suggested by equation 3A and the inducer's slant in the fashion suggested by equation 4. Slant contrast is predicted: the estimated slant of the test strip is offset in the direction opposite from the inducer's estimated slant. See text for further explanation.

The first term is a direct estimate that is based on signals created by the test strip alone and the second term is an indirect estimate based on signals created by the inducer alone and the relative disparity gradient between the inducer and test strip ($w_{t,dir} + w_{t,ind} = 1$).[8] If $w_{t,ind}$ is greater than 0 and $\hat{S}_{t,alone} \neq \hat{S}_{i,alone} + \hat{S}_{rel}$, the perceived slant of the test strip will differ from the direct slant estimate.

The inducer's slant estimate is in principle subject to the same constraints:

$$\hat{S}_i = w_{i,dir}\hat{S}_{i,alone} + w_{i,ind}(\hat{S}_{t,alone} - \hat{S}_{rel}), \qquad 3B)$$

but we assume that the influence of the test strip on the inducer's perceived slant is insignificant in the standard stimulus because the reliability of the inducer's slant estimate is much greater than that of the

[8] The model developed here assumes that the three disparity-based estimates (one from the test strip, one from the inducer, and one from the slant difference between them) are independent. This assumption is, however, false because the relative disparity gradient shares signals with the disparity gradients from the test strip and inducer. Thus, one should include an analysis of the covariation among these estimates. We have chosen to ignore this complexity in this paper and take it up in some forthcoming work. The effect of including covariation would change the weight estimates, but not the overall theoretical analysis presented here.

test strip's estimate (because the inducer is much larger and contains much richer perspective information).[9] Specifically, $w_{i,dir} \approx 1$ (and $w_{i,ind} \approx 0$), so

$$\hat{S}_i \approx \hat{S}_{i,alone}. \tag{4}$$

The perceived slant of the inducer thus depends only on the values and reliabilities of $\hat{S}_{id}$ and $\hat{S}_{ip}$ (equation (1A)).

As implied by equation (3A), the test strip's slant is estimated by combining the direct and indirect estimates. The relative reliabilities of the two estimates determine their weights. In general, the indirect estimate will be no more reliable than its least reliable term (e.g. $\hat{S}_{i,alone}$ and $\hat{S}_{rel}$ in equation 3A).

If $\hat{S}_{i,alone}$ and $\hat{S}_{rel}$ are both much more reliable than $\hat{S}_{t,alone}$, then their sum, though less reliable than either estimate alone, will also be more reliable than $\hat{S}_{t,alone}$. If we make this assumption, then $w_{t,ind} \approx 1$. By combination with equations (3A), (1A), and (2), and the assumption of unbiased estimators[10], we obtain:

$$S_t = S_{td} - w_{ip}(S_{id} - S_{ip}). \tag{5}$$

The predicted perceived slant of the test strip is thus determined by its physical slant (S_{td}), by the inducer's stereo- and nonstereo-specified slants (S_{id}, S_{ip}), and by the weight of the nonstereo-specified slant of the inducer (w_{ip}).

In most demonstrations of the slant-contrast effect, the test strip's stereo slant (S_{td}) and the inducer's nonstereo slant (S_{ip}) are zero. The predicted perceived slants of the inducer and test strip are then just:

$$\hat{S}_i = w_{id} S_{id} \tag{6A}$$
$$\hat{S}_t = -w_{ip} S_{id}. \tag{6B}$$

The predicted inducer slant is proportional to the stereo-specified slant (S_{id}) and to the weight given this slant. The predicted slant of the test strip is negatively proportional to the inducer's stereo-specified slant (S_{id}) and to the weight given the inducer's nonstereo-specified slant. The prediction of slant opposite in direction to the inducer's stereo slant is qualitatively consistent with observation. The bottom panel of Figure 2 displays predicted perceived slants for the test strip and inducer assuming that w_{id} is slightly smaller than w_{ip}.

The experiments reported here employed a slant-nulling procedure. Observers adjusted the physical slant of the test strip until it appeared to have zero slant: that is, until $\hat{S}_t = 0$. In this situation,

$$S_{td} = w_{ip}(S_{id} - S_{ip}). \tag{7A}$$

In the standard slant-contrast stimulus, $S_{ip} = 0$, so we obtain:

$$S_{td} = w_{ip} S_{id}. \tag{7B}$$

The nulling slant (S_{td}) should, therefore, be proportional to the inducer's stereo-specified slant (S_{id}) in the standard stimulus. We will return to this formulation when we present the data from Experiment 1.

---

[9] We assume here that the test strip has very little influence on the perceived slant of the inducer. This assumption is validated by some observations of van Ee (1995). Specifically, he showed that the strip can cause a local distortion of the inducer's perceived slant when the strip contains rich perspective information. Otherwise, the strip seems to have no influence on the inducer's perceived slant.

[10] We assume in the development of the equations that estimated slants (e.g., $\hat{S}_{id}$) are on average equal to their physical counterparts (e.g., S_{id}).

In the framework of slant estimation theory, the slant-contrast effect is a byproduct of the visual system's attempt to reconcile differing slant estimates that arise from the inconsistency between the stereo-based and nonstereo-based signals produced by the inducer. It follows then that slant contrast would not be observed if the inconsistency created by the inducer were eliminated by making the stereo and nonstereo signals specify the same slant for the inducer. With consistent signals ($S_{id} = S_{ip}$), the test strip's perceived slant ought to be (by equation (5)):

$$\hat{S}_t = S_{td}. \tag{8}$$

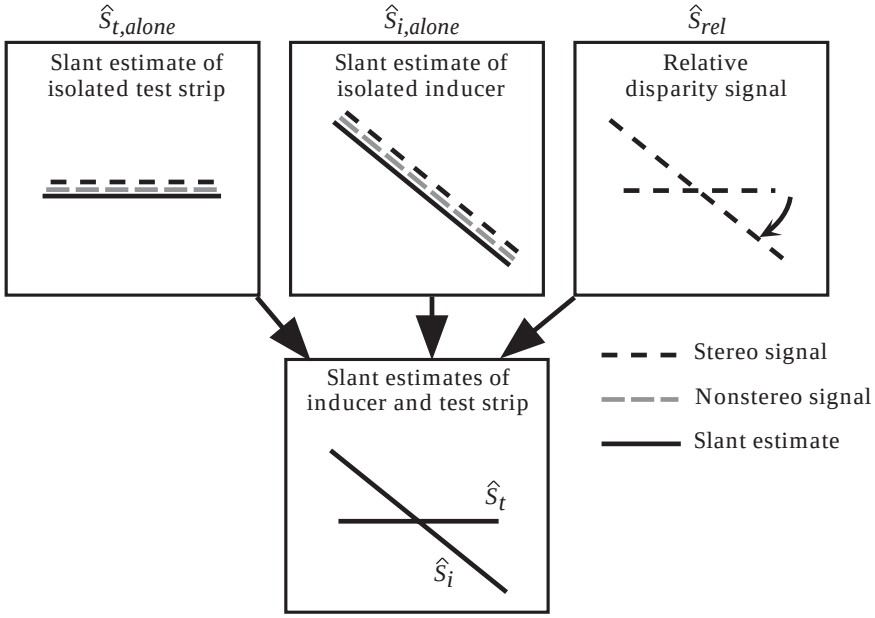

Figure 3. Slant estimation when the inducer's nonstereo-specified and stereo-specified slants are the same. Upper left panel: slant estimation for the isolated test strip. Upper middle panel: slant estimation for the isolated inducer. The stereo and nonstereo slant estimates are equal to one another, so according to the slant estimation theory presented in the text, the estimated slant of the inducer (solid black line) is the same as the slant indicated by either source alone. Upper right panel: estimation of the slant difference between the inducer and test strip from the relative disparity gradient. Bottom panel: final slant estimates for the test strip and inducer. The test strip's slant is estimated in the fashion suggested by equation 3A and the inducer's slant in the fashion suggested by equation 4. No slant contrast is predicted. See text for further explanation.

In other words, slant contrast should not be observed. Figure 3 schematizes this analysis when the inducer's stereo and nonstereo-specified slants are the same. When the task is slant nulling, equation (8) implies that the slant will be adjusted until $S_{td} = 0$, which means again that slant contrast should not be observed. The perceived slant of the inducer should be:

$$\hat{S}_i = S_{id},$$

so the inducer's perceived slant should not depend on the weights given to its stereo- or nonstereo-specified slants.

As developed above, the slant estimation model does not distinguish between slants about vertical and horizontal axes, so it predicts identical slant contrast effects in the two cases.

3 General Methods

In the experiments, we investigated the predictions of the slant estimation model. To do so, we modified the slant information in the inducer and measured the resulting effects on slant contrast.

3.1 Observers

Five observers with normal stereopsis participated. All were unaware of the experimental hypotheses. Refractive error was corrected by spectacles. All observers participated in a 40-trial training session in which they estimated the perceived slant of real and simulated planes; feedback was given during the training session. Three of the initial eight recruits were not able to make consistent slant estimations, so they did not participate in the main experiments.

3.2 Stimuli

The stimuli consisted of two planes that rotated about the same vertical axis (or horizontal axis). One plane–the inducer–had a regular, cross-hatch texture and a 3° high (or wide) gap extending across its full width (or height). The other plane–the test strip–was positioned in the gap (such that the two visible portions of the inducer flanked the test strip) and had a sparse, random-dot texture (Figure 1 is a simplified version). There were 0.5° gaps between the edges of the test strip and the edges of the inducer; these gaps were black. A fixation point was constantly visible at the midpoint of the test strip and observers were told to maintain fixation on that point.

In the vertical-axis condition, each inducer plane subtended 37.3 × 12.0° and the test strip 37.3 × 2.0°. In the horizontal-axis condition, we used the same configuration rotated through 90°, so the inducer planes subtended 12.0 × 37.3° and the test strip 2.0 × 37.3°.[11]

4 Experiment 1: Slant Contrast with Consistent and Inconsistent Texture Gradients

We measured slant contrast using standard and modified stimuli. In the standard stimulus, the nonstereo-specified slant of the inducer was zero. In the modified stimulus, the inducer's nonstereo-specified slant was the same as the stereo-specified slant. In both cases, the slant-contrast effect was measured with vertical- and horizontal-axis variations.

4.1 Apparatus and Stimuli

The stimuli were presented dichoptically using a conventional red-green anaglyphic technique. The stimuli were generated on a PowerMac 9500/132 using MATLAB and the OpenGL graphics libraries. They were projected by an Electrohome ECP-4000 projector onto a large flat screen. Frame rate was 75 Hz. At the viewing distance of 100 cm, each pixel subtended 4.3 minarc. The intensities of the red and green half-images were adjusted until they appeared equiluminant when viewed through the red and green filters. There was no visible crosstalk between the half-images. The room was very dark so that the screen edges and other environmental features could not be seen. The head was stabilized with a chin and forehead rest. Observers were told to maintain fixation on a central fixation point.

In the standard slant-contrast stimulus, the slant of the inducer is specified by disparity alone (Werner, 1938; Howard & Rogers, 1995). Specifically, the texture is mapped onto the two half-images and then the width of one half-image is increased in order to create the disparity gradient consistent with a plane rotated about a vertical axis. This manner of creating an apparent slant does not alter the texture gradient in the two half-images, so the texture gradient signals a slant of zero. Here we refer to this as the *inconsistent stimulus* because stereo and nonstereo cues signal different slants for the inducer. With the inconsistent stimulus, the texture of the inducer was a lattice of bright diagonal lines 8.6 minarc in width and separated by 3.5°; the inducer was otherwise black. The diagonal lines created square lattice elements on the screen (before the disparity gradient was modified). When the disparity gradient was modified by stretching one half-image horizontally (and compressing the other), the lattice elements became rhombuses on the screen; the texture gradient remained that of a frontal plane.

The test strip was textured with approximately 20 dots (diameter = 8.6 minarc) at random positions.[12] The positions of the dots were randomly reassigned on each stimulus presentation.

We also created a version of the slant-contrast stimulus in which the stereo-specified and nonstereo-specified slants of the inducer were identical. We call this the *consistent stimulus* because stereo and

[11] In the consistent condition in which the disparity- and perspective-specified slants of the inducer were the same, the projected width of the stimulus naturally varied with its slant.

[12] In order to determine the informativeness of any residual texture-gradient cues, we performed a control experiment. Observer JVE viewed the test strip monocularly and set its apparent slant to gaze normal. The standard deviation of her settings was 7.4 deg, which is 5–10 times greater than the standard deviations of her binocular settings. Thus, residual texture-gradient cues in the test strip were much less informative than the disparity cue.

nonstereo cues signal the same inducer slant. For this stimulus, the projection of the texture to the two eyes was calculated with the OpenGL libraries to be appropriate for the disparity gradient and distance of the inducer. The orientations of and separations between the texture lines were geometrically correct for a slanted surface viewed at the appropriate station point and so were the line thicknesses (to within the precision of our anti-aliasing algorithm). The outline shape of the inducer was also geometrically correct for a slanted rectangle. The test strip was identical in the inconsistent and consistent stimuli; only the inducer was affected by this manipulation.

Recall that the normalization and inhibition models predict that slant contrast will either increase or remain constant when the nonstereo-specified slant of the inducer is made consistent with its stereo-specified slant. Our model predicts that slant contrast will be eliminated by this manipulation (equation (8)).

The stereo-specified slant of the inducer (vertical- and horizontal-axis rotations) was –30, –15, 0, 15, or 30˚ in both the consistent and inconsistent configurations. The simulated distance to the midpoint of the inducer (and test strip) was 38, 190, or 570 cm; the distance was specified by the disparity of the fixation point and by the gradients of horizontal and vertical disparity at the retinae. We included a manipulation of simulated distance because of the expectation that nonstereo slant cues should be given greater weight with increasing distance (Sedgwick, 1986; Backus & Banks, 1999; Johnston, Cumming, & Parker, 1993). The disparity manipulations were done for an assumed interocular distance of 6.4 cm, which is a representative value for the five observers. The projected size of the inducer and test strip did not vary with simulated viewing distance.

4.2 Procedure and Task

Observers initiated stimulus presentations with a button press and then the inducer and the test strip appeared simultaneously for 1.5 sec. The initial, randomized slant of the test strip (specified stereoscopically) was –10 to 10˚. Observers adjusted the strip's slant until it appeared "parallel to the face"; that is, they nulled its perceived slant. Adjustments were made after the stimulus had been extinguished. After each adjustment, the stimulus reappeared with the test strip rotated to a new slant. Adjustments continued until the observer was satisfied. Feedback was not provided. Between stimulus presentations, a frontoparallel test strip appeared in order to remind the observer of the criterion; in a control experiment with a subset of the observers, we found that removal of the reminder stimulus did not affect the average slant settings, but it did increase the variability.

After observers completed a test strip setting, the task changed to indicating the inducer's perceived slant. The same stimulus (with the test strip in the "nulled" position) was presented again for 1.5 sec followed by two intersecting line segments. The orientation of one of the segments was fixed (horizontal in the vertical-axis condition and vertical in the horizontal-axis condition) and the orientation of the other could be adjusted by moving the computer mouse. The fixed line segment represented the frontoparallel plane, so observers adjusted the orientation of the other segment until it indicated the inducer's perceived slant relative to the frontoparallel plane (van Ee & Erkelens, 1996a).

Each observer was tested in six sessions (3 distances with 2 slant-axis conditions). The first three sessions involved vertical-axis slants and the last three horizontal-axis slants. Each session consisted of 50 trials presented in random order: two nonstereo conditions (consistent and inconsistent), five inducer slants (-30, -15, 0, 15, and 30 deg), and five repetitions per condition.

4.3 Results

The results are displayed in Figures 4 and 5, the former for the vertical-axis condition and the latter for the horizontal-axis. In both cases, the upper panels plot the estimated slant of the inducer as a function of its stereo-specified slant and the lower panels plot the slant of the test strip when it appeared parallel to the face. The left, middle, and right panels show the data for simulated distances of 38, 190, and 570 cm, respectively. The filled and unfilled symbols represent data from the consistent and inconsistent conditions, respectively. Each data point represents the slant settings averaged across the five observers.[13]

[13] The data from the five observers were very similar to one another, so the data averaged across observers is quite representative of individual behavior. Occasionally, two observers reported inducer slants in the incorrect direction (slant reversals; Gillam, 1967). The test strip and inducer data from those trials were discarded and thus are not represented in the averages.

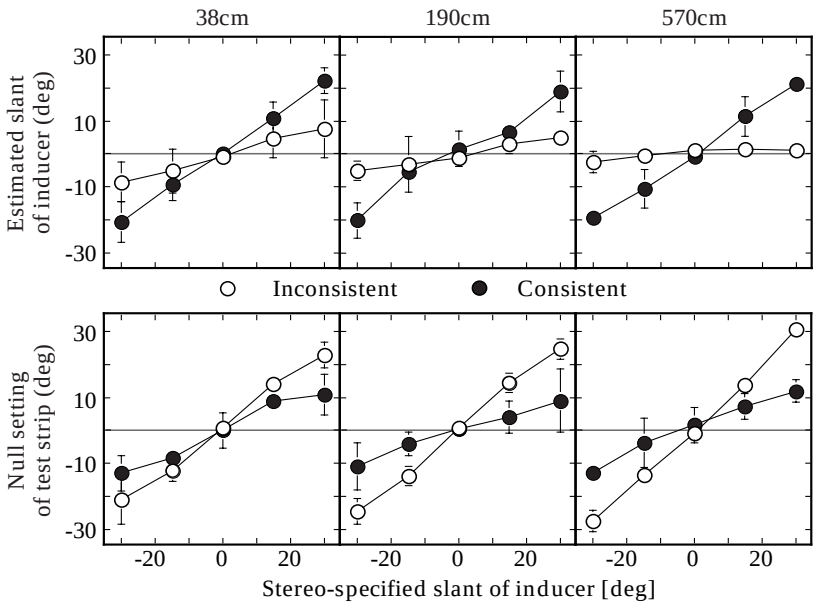

Figure 4. Results for vertical-axis condition in Experiment 1. The left, middle, and right columns show data for simulated distances of 38, 190, and 570 cm, respectively. Unfilled symbols represent data for the inconsistent condition (inducer's stereo- and nonstereo-specified unequal) and filled symbols data for the consistent condition (inducer's stereo- and nonstereo-specified equal). Data points are the average of five observers' data. Error bars represents +/-1 standard deviation of the five averages. Upper panels: Estimated slant of the inducer as a function of its stereo-specified slant. Lower panels: Objective slant of the test strip when it appeared frontoparallel as a function of the inducer's stereo-specified slant.

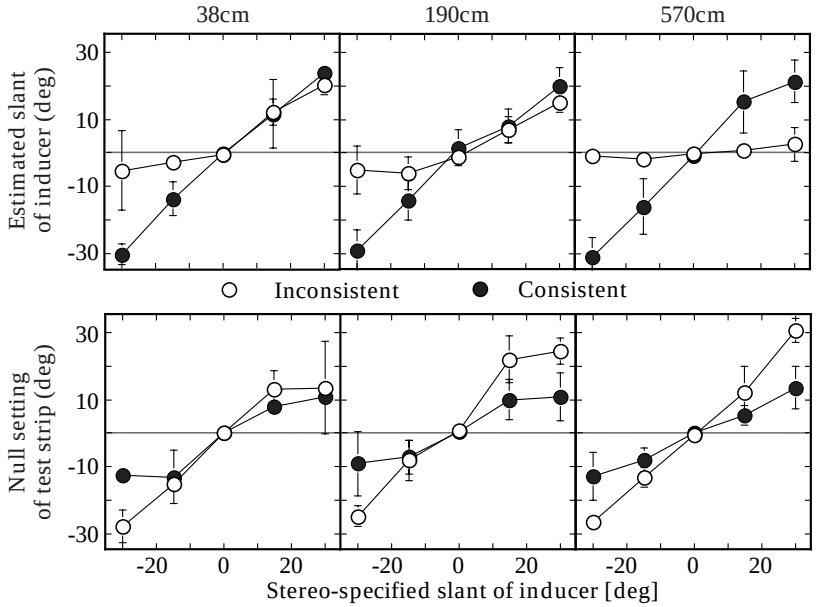

Figure 5. Results for horizontal-axis condition in Experiment 1. Conventions the same as in Figure 4.

First consider the perceived slants reported for the inducer (upper panels). In the inconsistent condition (unfilled symbols), observers reported much lesser slants than the actual slants specified stereoscopically, particularly at the longest viewing distance. From Figures 4 and 5 and equations (1A) and (4), we can estimate the weight (w_{ip}) given to the inducer's nonstereo-specified slant when in conflict with its stereo-specified slant. To make such estimates, one has to assume that the function mapping perceived slant onto the observer's response (setting one line's orientation relative to another) is veridical, an assumption that is probably not valid. If we assume, however, that any error in this mapping is consistent across conditions of the experiment, we can compare how the weight estimates

change with experimental condition. The estimates of w_{ip} are provided in Table 1. The estimated nonstereo weights increase systematically with viewing distance (and the stereo weights decrease) as expected from an analysis of signal reliability (Backus & Banks, 1999). There is a tendency for the estimated nonstereo weight to be higher for the vertical- than for the horizontal-axis condition which is consistent with previous reports concerning slant anisotropy (Mitchison & McKee, 1990; Gillam & Ryan, 1992; Buckley & Frisby, 1993).

	Distance (cm)		
	38	190	570
Vertical-axis	0.72	0.83	0.94
Horizontal-axis	0.56	0.65	0.94

Table1. Weights (w_{ip}) Estimated from Inducer's Reported Slants (Inconsistent Condition, Figures 4 and 5) (These values were estimated from the slant estimates in the upper panels of Figures 4 and 5 using the equation: $w_{ip} = 1 - \hat{S}_i / S_{id}$.)

The inducer's reported slants were systematically greater in the consistent condition (filled symbols) than in the inconsistent condition. Thus, as expected, when the stereo- and nonstereo-specified slants are the same, reported slants are closer to the depicted values. Notice, however, that the reported slants in the consistent condition were still lower than the depicted values. This shortfall could be a consequence of a non-veridical mapping between perceived slant and the observers' responses or to the presence of other cues (such as those provided by the display screen) indicating that the simulated plane is frontoparallel; we will consider the latter possibility in Experiment 2.

Now consider the test strip data (lower panels). If there were no effect of the inducer on the test strip's perceived slant, observers would set the strip's slant to zero for all conditions. If, on the other hand, the inducer produced slant contrast, observers would set the test strip's slant in the direction of the inducer's stereo slant (equation (7B)). This happened for all conditions. The most important observation, however, is that slant contrast was reduced by making the nonstereo-specified slant of the inducer consistent with its stereo-specified slant; in all six conditions (three in each of Figures 4 and 5), the null settings were closer to zero in the consistent than in the inconsistent condition. This effect is somewhat clearer at the longest viewing distance.

From equation (7B), the slopes of the null-setting data in the inconsistent condition provide another estimate of w_{ip}. Those estimates are provided in Table 2. The estimates of w_{ip} in Tables 1 and 2 are similar which confirms our assumption that the indirect term in equation (3A) is given very high weight. By implication, the objective slant of the test strip has little, if any, effect on its perceived slant apart from its contribution to the relative disparity gradient.

Table 1 shows that we observed some vertical-horizontal differences in the nonstereo weight (w_{ip}) in the estimates of the inducer's slant. Table 2 shows that we did not observe such differences in the null settings. We do not have an explanation for this slight discrepancy between the two means of estimating w_{ip}.

	Distance (cm)		
	38	190	570
Vertical-axis	0.76	0.84	0.96
Horizontal-axis	0.73	0.86	0.93

Table 2. Weights (w_{ip}) Estimated from Slopes of Nulling Data (Inconsistent Condition, Figures 4 and 5) (These values were estimated from the slant estimates in the upper panels of Figures 4 and 5 using the equation: $w_{ip} = S_{td} / S_{id}$.)

Previous models predict that slant contrast will either increase or remain the same when one changes the inducer's nonstereo-specified slant from zero to the value that is consistent with the stereo-specified slant. Our data are clearly inconsistent with this prediction because in every case slant contrast decreased with this manipulation.[14] Slant estimation theory predicts that slant contrast will be reduced when the inducer's slant is specified in a consistent fashion and that prediction is confirmed by the

[14] This empirical observation has also been reported recently by Sato and Howard (1999).

data. However, the theory actually predicts that *no* slant contrast should be observed in the consistent condition and our data do not confirm this prediction.

Why might slant contrast still be observed when the stereo- and nonstereo-specified slants of the inducer are consistent? Images on computer display screens provide a number of signals that specify the slant of the screen and not the slant of simulated surfaces. For example, when the observer fixates along a normal to a flat display screen, the change in blur with azimuth is zero or nearly zero because the vergence of the incoming light is determined by the distance from the eye to the screen; this cue specifies a slant of 0˚ no matter what object slant is actually being simulated. Additionally, at the 100-cm viewing distance used here, individual pixels can be resolved and the graininess of the simulated surface also specifies a slant of 0˚ rather than the simulated value.

We can incorporate this class of slant signals in our theory by splitting the nonstereo slant estimators, $\hat{S}_{ip}$ and $\hat{S}_{tp}$, into two estimators each, one based on nonstereo signals that can be rendered correctly on the display screen and the other based on nonstereo signals that cannot be rendered correctly[15]:

$$\begin{aligned} w_{ip}\hat{S}_{ip} &= w_{ic}\hat{S}_{ic} + w_{iu}\hat{S}_{iu} \\ w_{tp}\hat{S}_{tp} &= w_{tc}\hat{S}_{tc} + w_{tu}\hat{S}_{tu} \end{aligned} \tag{9}$$

where p refers to all nonstereo signals, c to controllable signals, and u to uncontrollable signals. Furthermore, $w_{ic} + w_{iu} = w_{ip}$ and $w_{tc} + w_{tu} = w_{tp}$. The equations for slant estimation with the isolated surfaces (from equation (1)) become:

$$\begin{aligned} \hat{S}_{i,alone} &= w_{id}\hat{S}_{id} + w_{ic}\hat{S}_{ic} + w_{iu}\hat{S}_{iu} \\ \hat{S}_{t,alone} &= w_{td}\hat{S}_{td} + w_{tc}\hat{S}_{tc} + w_{tu}\hat{S}_{tu} \end{aligned} \tag{10}$$

We presume that $\hat{S}_{iu}$, the slant estimator based on uncontrollable nonstereo signals, always indicates a frontoparallel surface in our experimental situation, so provided that the weight w_{iu} is non-zero, that signal will attenuate the inducer's perceived slant. From the earlier discussion, we know that signals causing attenuation of the inducer's perceived slant relative to its stereo-specified slant are crucial to the creation of the slant-contrast effect. Thus, we hypothesize that the residual slant contrast observed in the consistent condition of Experiment 1 was caused by the presence of the uncontrollable nonstereo signals. We tested that hypothesis in the next experiment.

5 Experiment 2: Slant Contrast with Real Planes

We wished to see if slant contrast disappears, as predicted by slant estimation theory, when *all* signals provided by the inducer specify the same slant. We did so by re-doing the experiment with real objects rather than with computer graphic images.

5.1 Apparatus and Stimuli

We mimicked the methodology of Experiment 1 to the extent possible. The same five observers participated. The inducer and test strip were constructed from real planes attached to a common rotation axis. The two parts of the inducer were connected by an invisible support structure, so they were always co-planar. Both inducer planes were 70 x 20 cm and the test strip was 70 x 3.5 cm; at the 100-cm viewing distance, they subtended the same angles as the images in Experiment 1. The texture of the inducer was a black-and-white checkerboard, each check subtending 1.5˚. The inducer's slant was varied by the experimenter and measured using a protractor and sighting device. The test strip was flat black with sparse white dots added in random positions; this texture had roughly the same properties as the test strip's texture in the first experiment. A central dot served as the fixation point. Observers adjusted the test strip's slant with a rope-and-pulley system; the system was arranged such that rope and hand position did not provide reliable cues to objective slant.

[15] We posit two separate estimators for purposes of analyzing the data. We do not claim that the visual system actually processes the controllable and uncontrollable signals separately.

The illumination of the apparatus was carefully designed to eliminate artifactual slant cues such as shadows. Observers viewed the stimulus through a large, half-silvered mirror that was placed in front of the face and rotated 45° to the line of sight. The illumination source was a small tungsten bulb at eye level. Light from the source reflected off the mirror and illuminated the inducer and test strip. The optical path length was 100 cm, so the source was conjugate with the observer's cyclopean eye. The apparatus was placed in front of an invisible black curtain and thus relative disparities between apparatus features and the background could not be measured. Other features in the room were rendered invisible with curtains and paint. The room lights were turned on between trials to insure that observers did not dark-adapt to the point where they could see extraneous features. Observers' head position was again stabilized with a chin and forehead rest. We instructed observers to maintain completely fixed head position while viewing the stimulus.

The whole apparatus (except the light source and mirror) was rotated by 90° depending on whether the vertical- or horizontal-axis condition was being presented.

5.2 Procedure

The experimenter set the inducer's slant on each trial to a value of –30, –15, 0, 15 or 30°. He also set the test strip's slant to a random value of –6 to 6°. Observers held the eyes closed during these initial settings.

Observers adjusted the slant of the test strip until it appeared to be parallel to the face. Adjustments were made with discrete movements of the rope-and-pulley system. The eyes were closed during the adjustments. Observers then opened the eyes to re-assess the apparent slant, closed the eyes, made another adjustment, and so forth until they were satisfied. (The eyes were closed during the actual adjustments so that observers could not use motion parallax to judge the strip's slant.) After each setting, observers indicated the apparent slant of the inducer using the same procedure as in Experiment 1. The line segments were presented on a display screen placed on the opposite side of the room from the light source; observers had to turn the head to do that task.

Each observer made a total of 50 test-strip settings and 50 inducer slant estimates (2 slant axes, 5 inducer slants, and 5 repetitions of each condition).

5.3 Results

The results are shown in Figure 6 in the same format as the previous data figures. The unfilled and filled symbols now represent the vertical-axis and horizontal-axis settings, respectively. The upper panel shows the slant estimates for the inducer, averaged across observers; the estimates were quite accurate. The lower panel shows the test strip's objective slant when it appeared parallel to the face. Remarkably, the null settings were close to 0° for all conditions. As predicted by the theory, slant contrast was essentially abolished when all slant signals were made consistent by using real objects[16]. This finding further supports the idea that the illusory slant seen in the standard slant-contrast stimulus is caused by conflicts between various slant signals associated with the inducer. Observers reported greater perceived slants for the inducer in this experiment than they did in Experiment 1. Previous theories predict that slant contrast should either increase or remain the same when the perceived slant of the inducer increases, but we observed a decrease in slant contrast.

By comparing the results of Experiments 1 and 2, we can estimate the weight given to the uncontrollable nonstereo signals created by projection onto a flat display screen. In the consistent condition of Experiment 1, the uncontrollable signals specified a frontoparallel plane, so using the reasoning that led to equation (7A):

$$S_{td} = w_{iu} S_{id} . \qquad (11)$$

[16] Experiment 2 was conducted with binocular viewing, but it is important to consider the possibility that the null settings were based on nonstereo slant signals from the test strip alone. To examine this possibility, we conducted a monocular control experiment with three of the five observers. They performed the same task as in Experiment 2 except they viewed with the right eye only. The slant estimates for the inducer were similar to those shown in the top panel of Figure 6. The standard deviations of the null settings increased by factors of 2.8, 3.5, and 4.7 compared with the binocular settings, but were otherwise similar. This suggests that observers relied on binocular signals while doing the null settings in the main experiment.

If w_{iu}, the weight given uncontrollable nonstereo signals, is greater than 0, then we expect the null setting (S_{td}) to be proportional to the inducer's slant (S_{id}) in the consistent condition of Experiment 1 (Figures 4 and 5). The weight w_{iu} can be estimated from the null-setting slopes in Experiment 1. The weights estimated in this fashion are given in Table 3; the average value is 0.40 (that is, 40% of the total weight given to nonstereo), which suggests that uncontrollable nonstereo signals had a significant effect on the data in Experiment 1.

	Distance (cm)		
	38	190	570
Vertical-axis	0.43	0.32	0.40
Horizontal-axis	0.45	0.38	0.44

Table 3. Weights for Uncontrollable Nonstereo Signals (w_{iu}) Estimated from Slopes of Nulling Data (Consistent Condition, Figures 4 and 5) (These values were estimated from the slant estimates in the upper panels of Figures 4 and 5 using the equation: $w_{iu} = S_{td} / S_{id}$.)

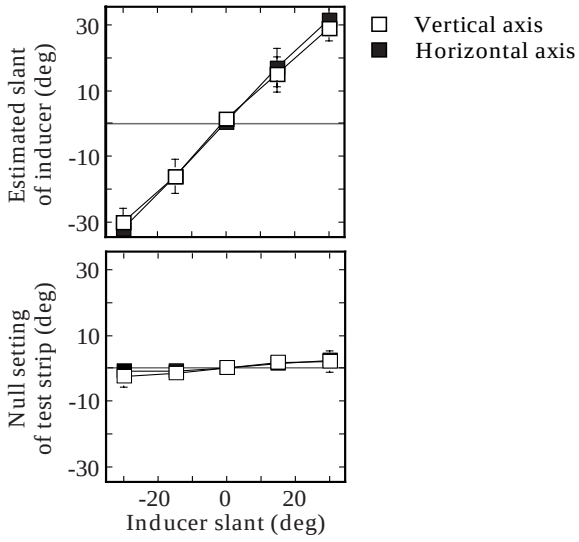

Figure 6. Results for Experiment 2 when the stimulus consisted of real planes. Unfilled symbols represent data for the vertical-axis condition and filled symbols data for the horizontal-axis condition. Upper panels: Estimated slant of the inducer as a function of its objective slant. Lower panels: Objective slant of the test strip when it appeared frontoparallel.

6 Experiment 3: Reverse Slant Contrast

According to slant estimation theory, the illusory slant seen in the standard slant-contrast stimulus is the byproduct of inconsistencies between slant signals from the inducer. The theory makes an interesting prediction when the signal inconsistency from the inducer is reversed: it predicts slant contrast in the opposite direction when the inducer's stereo-specified slant is zero and its nonstereo-specified slant is non-zero.

The signals and prediction in this situation are depicted in Figure 7. One can understand the prediction from consideration of equation (5). The inducer's stereo-specified slant (S_{id}) is zero and its nonstereo-specified slant (S_{ip}) is non-zero, so equation (5) becomes:

$$\hat{S}_t = S_{td} + w_{ip} S_{ip}.$$

For our nulling task ($\hat{S}_t = 0$), we obtain:

$$S_{td} = -w_{ip} S_{ip}.$$

The null setting of the test strip should be opposite to the direction of the inducer's perceived slant, which is *opposite* to the direction of the effect in the standard stimulus.

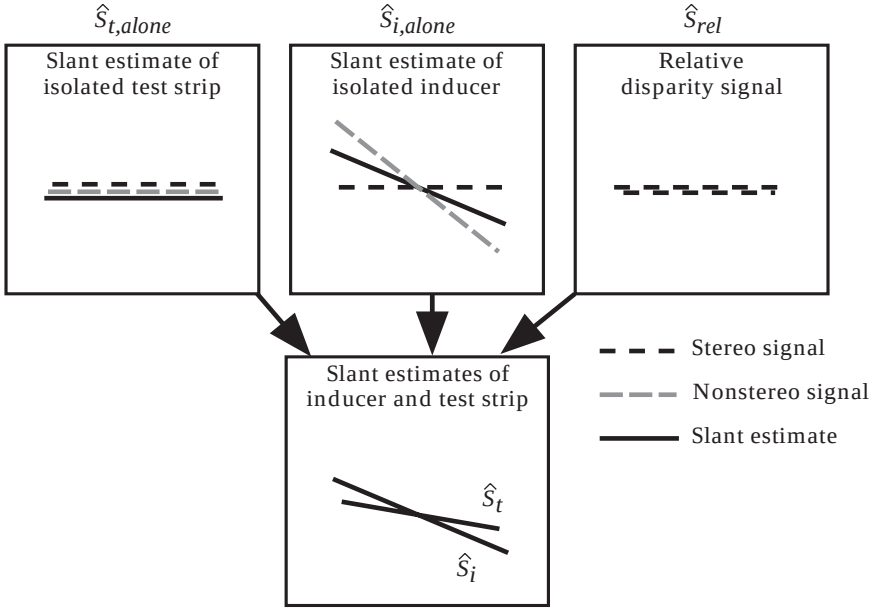

Figure 7. Slant estimation when the inducer's nonstereo-specified slant is nonzero and its stereo-specified slant is zero. The nonstereo- and stereo-specified slants are inconsistent as in the standard slant-contrast stimulus, but stereo signals indicate a frontal plane and nonstereo signals do not (the reverse of the situation in Figures 1 and 2). Upper left panel: slant estimation for the isolated test strip. Upper middle panel: slant estimation for the isolated inducer. Upper right panel: estimation of the slant difference between the inducer and test strip from the relative disparity gradient between the two surfaces. The stereo-specified slants of the inducer and test strip are the same, so the relative disparity gradient between the two is zero indicating no slant difference. Bottom panel: final slant estimates for the test strip and inducer. Reverse slant contrast is predicted. See text for further explanation.

6.1 Apparatus and Stimuli

The observers, stimuli, and procedure were the same as in Experiment 1 with the following exceptions. 1) The inducer's stereo-specified slant was always 0°, so its perceived slant was varied by altering its nonstereo-specified slant. 2) Only two observers participated. 3) Only one simulated viewing distance was presented: 570 cm. 4) Slant was varied about the vertical axis only.

Again observers adjusted the slant of the test strip until it appeared parallel to the face; they also estimated the inducer's slant as in the previous experiments.

6.2 Results

The results are shown in Figure 8. The upper panel shows the slant estimates for the inducer, averaged across observers, as a function of the inducer's nonstereo-specified slant. The lower panel shows the objective slant of the test strip when it appeared parallel to the face. Notice that the strip's objective slant was always in the opposite direction from the nonstereo-specified slant of the inducer. In other words, reversed slant contrast was observed as predicted by slant estimation theory.

Controllable and uncontrollable nonstereo signals specify different slants in this experiment (the latter always specifying a slant of 0). We must take that fact into consideration when estimating weights from the data. We can estimate the weight w_{ic} from the slant estimation data in the upper panel (based on equation (1A) with appropriate modifications): the average estimated w_{ic} is 0.58. We can also estimate w_{ic} from the slope of the nulling data in the lower panel: the average estimated w_{ic} is 0.52.[17] The two estimates are very similar which confirms our earlier assumption that the indirect term in equation (3A) is given very high weight.

[17] We used the following two equations to estimate w_{ic}. From the estimated inducer slants: $w_{ic} = S_i / S_{ic}$. From the null settings of the test strip: $w_{ic} = -S_{td} / S_{ic}$.

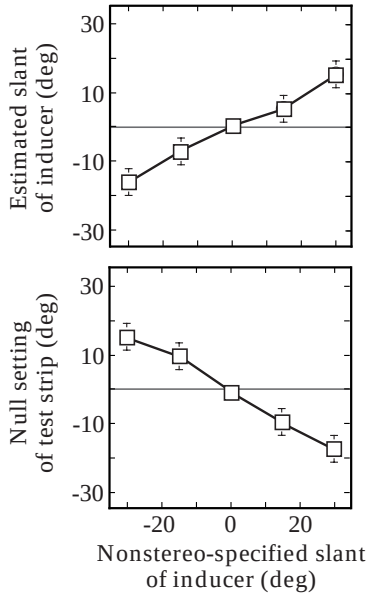

Figure 8. Results for Experiment 3 in which the inducer's nonstereo-specified slant is nonzero and its stereo-specified slant is zero. Data points are the average of two observers' data. Error bars represent +/-1 standard deviation of the two averages. Upper panel: Estimated slant of the inducer as a function of its nonstereo-specified slant. Lower panels: Objective slant of the test strip when it appeared frontoparallel.

7 Discussion

Our results demonstrate that the phenomenon of slant contrast depends strongly on conflicts between the inducer's slant signals. To illustrate this point, Figure 9 plots the magnitude of slant contrast (quantified by the observers' null settings divided by the slant specified in the inducer) for the various experimental conditions. The upper and lower panels show data from the vertical- and horizontal-axis conditions, respectively. The conflict between the inducer's stereo- and nonstereo-specified slants varies from left to right. In the inconsistent condition, stereo-specified slant varied, but controllable and uncontrollable nonstereo slants were always zero. In the consistent condition, stereo and controllable nonstereo slants varied together, but uncontrollable slant was zero. In the real-planes condition, all three slant signals varied together. Finally, in the reverse condition, controllable nonstereo slant varied, but stereo and uncontrollable nonstereo were zero. Figure 9 makes clear that this variation in conflict between the inducer's slant specifications has a systematic and profound effect on the magnitude and even the sign of slant contrast.

Two features are required for slant contrast: 1) inconsistency between the slant signals provided by the inducing background and 2) a non-zero relative disparity gradient between the test strip and the inducer.

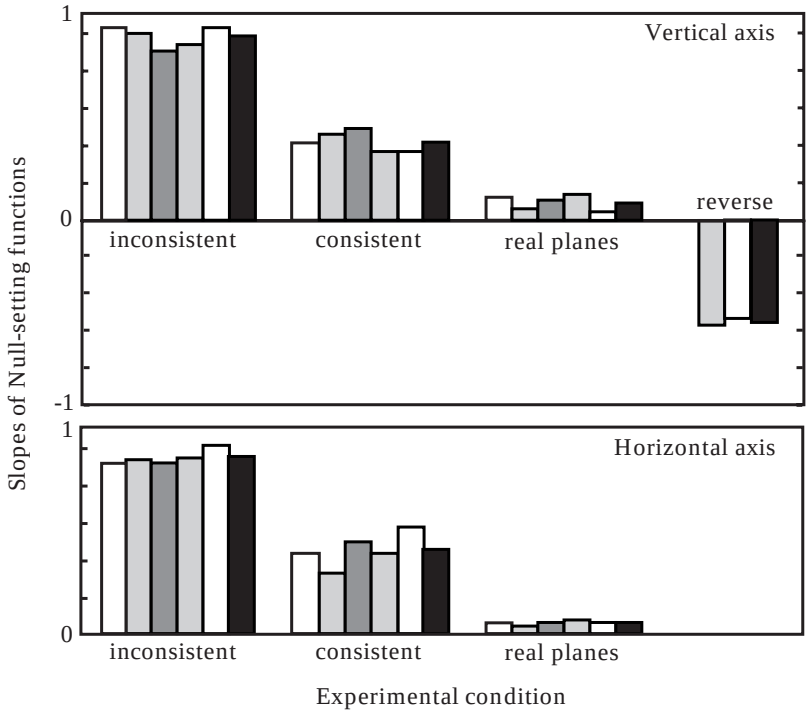

Figure 9. Summary of results for Experiments 1, 2, and 3. The slopes of the null-setting functions (objective slant of test strip when apparently frontoparallel divided by inducer slant) are plotted for individual observers for the 570-cm viewing distance. The upper graph displays data for vertical-axis rotations and the lower graph data for horizontal-axis rotations. "Inconsistent" and "consistent" refer to data from Experiment 1, "Real plane" to data from Experiment 2, and "Reverse" to data from Experiment 3. Positive slopes indicate that slant contrast was observed and negative slopes that reversed slant contrast was observed. In the "inconsistent", "consistent" and "real plane" data, the first five bars represent individual data and the sixth bar is the group average. In the "reverse" data, the first two bars are individual data and the third is the group average.

7.1 Comparison with Previous Results

Two reports have also claimed that inconsistencies between the slant signals provided by the inducer are essential to the phenomenon of slant contrast.

Ogle (1946) presented a stimulus in which a horizontal line (equivalent to the test strip) was surrounded by a rectangular frame (equivalent to the inducer). This stimulus is like the one in Figure 12.4(b) of Howard and Rogers (1995). Ogle claimed that he did not observe slant contrast with this stimulus because his observers were unable to fuse it. To facilitate fusion, he added perspective cues to the slant of the rectangular frame by making its outline shape trapezoidal. No details were provided, so it unknown whether the outline shape specified the same slant as specified stereoscopically. Ogle (1946) stated, without providing data, that slant contrast did not occur when the frame's outline shape was trapezoidal. This condition is similar in concept to the consistent condition of Experiment 1.

Sato and Howard (1999) measured slant contrast in conditions similar to the inconsistent and consistent conditions of Experiment 1. They too found that slant contrast is reduced, but not eliminated, when the inducer's nonstereo-specified is made consistent with stereo-specified slant.

7.2 Why is the Relative Disparity Gradient Given High Weight?

The slant estimation theory presented here states that the perceived slant of the test strip is $\hat{S}_t = w_{t,dir}\hat{S}_{t,alone} + w_{t,ind}(\hat{S}_{i,alone} + \hat{S}_{rel})$ (equation (3A)). One of the most striking findings is that the weight given the second term is very much higher than that given the first term; that is, the test strip's perceived slant seems to be determined more by its slant relative to the inducer (the quantity within the parentheses) than by signals coming from the test strip itself. Why does this occur? The reliability of each slant estimator depends on the reliability with which the signals it uses can be measured. $\hat{S}_{t,alone}$ can be known only from the texture and disparity signals within the test strip itself (equation (1B)). The

texture signal will be quite unreliable because the texture is sparse and random. Concerning the disparity signals, (1) why is the stereo estimate of the strip's slant (alone) much less reliable than that of the inducer's slant (alone), and (2) why is it much less reliable than the estimate of relative slant from stereo?

The answer to the first question is straightforward. According to slant estimation theory, the weight given a signal is proportional to the signal's reliability. The inducer is much larger than the test strip, so its stereo slant signal will be more reliable, and its regular texture provides an additional slant signal that is not available in the test strip. Thus, the inducer's slant signals are given higher weight because of the way the slant-contrast stimulus is constructed.

The answer to the second question involves consideration of how measurement errors affect the reliabilities of estimates made using horizontal disparity gradients and relative horizontal disparity gradients. The horizontal disparity gradient is ambiguous by itself. For example, the gradient created by a frontal plane presented in the head's median plane is identical to the gradient of any plane tangent to the Vieth-Müller circle (Backus et al, 1999; Ebenholtz & Paap, 1973; Gillam & Lawergren, 1983; Mitchison & Westheimer, 1990; van Ee & Erkelens, 1996c).[18] In order to interpret the horizontal disparity gradient, the visual system must compensate for the position of the surface patch relative to the head. Vertical disparity and extra-retinal, eye-position signals allow such compensation (e.g., Mayhew & Longuet-Higgins, 1982; Porrill, Mayhew & Frisby, 1990; Rogers & Bradshaw, 1995; Banks & Backus, 1998; Erkelens & van Ee, 1998; Backus et al., 1999).

It is instructive to compare how error in compensation affects the interpretation of the horizontal disparity gradient and the relative horizontal disparity gradient. Consider the two situations depicted in Figure 10. The left panel shows three surface patches, all tangent to the Vieth-Müller circle. The black patch is a frontal plane positioned straight ahead; it gives rise to a disparity gradient of zero. Surfaces at other azimuths create the same gradient if their slant is equal to their azimuth (Howard & Rogers, 1995). The gray diagonal line in Figure 11 represents those slants as a function of azimuth. Thus, a given error in compensation for azimuth (e.g., a 5˚ error in the estimated version of the eyes) yields an error of equal magnitude in estimating the slant of a surface patch (e.g., a 5˚ slant estimation error). The right panel in Figure 10 shows three pairs of intersecting surface patches; one member of each pair is tangent to the Vieth-Müller circle. The black pair is a frontal plane and a plane with a slant of –30˚; this pair gives rise to a particular relative disparity gradient. The gray pairs are positioned at other azimuths such that one patch is again tangent to the Vieth-Müller circle. The solid black curve at the bottom of Figure 11 shows the slant difference in those pairs that would yield the same relative disparity gradient as the black pair. The other curves represent the slant difference that would give rise to the same relative gradient as the black pair when its slant difference takes on different values. The slant difference is much less dependent on azimuth than the slant for a single surface is. Consider, for example, a slant difference of –30˚ (bottom curve): a 5˚ error in the estimated version (-2.5 vs 2.5˚) yields only a 1˚ error in the estimated slant difference. Thus, the relative horizontal disparity gradient is less sensitive to error in azimuth than is the horizontal disparity gradient.[19] For this reason, the reliability of absolute slant estimated from the horizontal disparity gradient ought to be lower than the reliability of the relative slant estimated from the relative gradient. Stated another way, the visual system probably knows the slant of the test strip relative to the inducer from the relative horizontal disparity gradient better than it knows the slant of the test strip alone from the horizontal disparity gradient.

[18] In other words, the same horizontal disparity gradient is created by a plane positioned straight ahead that has a slant of 0˚ and by a plane positioned 20˚ to the left of straight ahead that has a slant of 20˚.

[19] Error in distance estimation (e.g., vergence error) affects the interpretation of the horizontal disparity gradient and the relative horizontal disparity gradient, depending somewhat on the viewing parameters. Both types of slant error scale approximately linearly with distance, if azimuth is held constant. We believe, therefore, that version (or azimuth) error, not vergence (or distance) error, is the cause of the low weight given the disparity gradients of the isolated test strip and inducer and of the high weight given the relative disparity gradient between the strip and inducer.

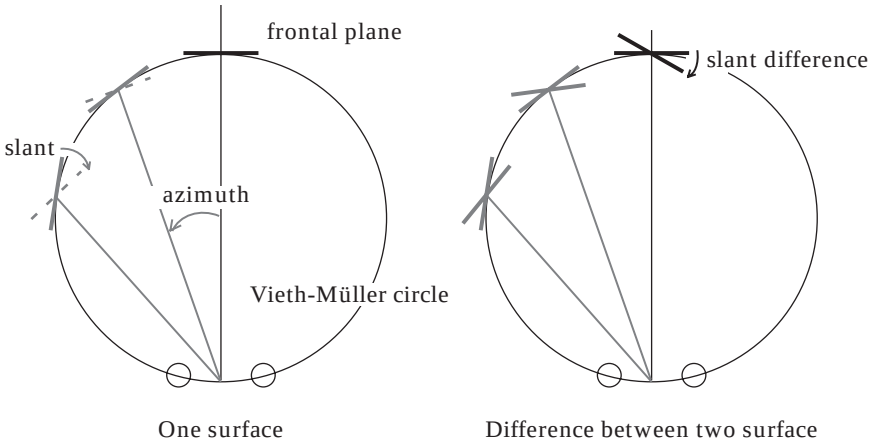

Figure 10. Estimating slant from the horizontal disparity gradient and slant difference from the relative horizontal disparity gradient. The left panel depicts a viewing situation in which the observer fixates a frontal plane straight ahead (black line segment); the plane creates a horizontal disparity gradient of zero locally. Planes at other azimuths, but tangent to the Vieth-Müller circle, create the same disparity gradient (gray segments). The right panel depicts a viewing situation in which the observer fixates two intersecting planes (black pair) straight ahead; the planes differ in slant by the indicated angle. That pair of planes creates a specific relative horizontal disparity gradient. Pairs at other azimuths can create the same relative disparity gradient, although the slant difference will have to be slightly different in order to do so. The pairs shown have one segment normal to the line of sight and the other at a different slant.

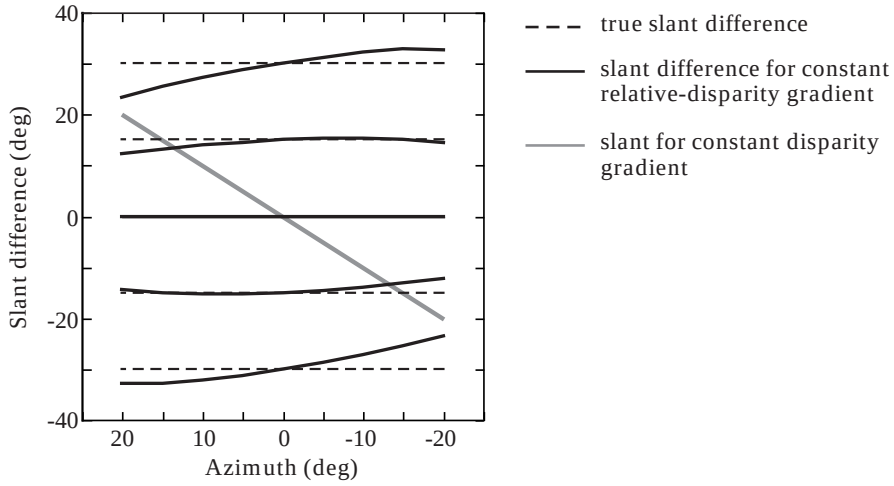

Figure 11. Uncertainty in slant estimation from the horizontal disparity gradient and in slant difference estimation from the relative horizontal disparity gradient. The gray line shows the actual slant of a patch constrained to lie on the Vieth-Müller circle that gives rise to the same horizontal disparity gradient as a frontal patch that lies straight ahead (see left panel, Figure 10). Slant is equal to azimuth, so a 5% error in azimuth compensation will yield a 5% error in the estimated slant. The black curves show the actual slant difference of a patch pair that gives rise to the same relative horizontal disparity gradient as a pair that lies straight ahead (right panel in Figure 10). For these calculations, the intersection of the pair straight ahead was at a distance of 38 cm. Interocular distance was 6.4 cm. One member of the pair was normal to the line of sight and the other member took on slant values (from bottom to top) of –30, –15, 0, 15, and 30˚ (dashed horizontal lines). We then considered pairs at other azimuths; their intersection lay on the Vieth-Müller circle and one member was normal to the line of sight. We calculated the slant difference that would create the same relative disparity gradient. The black curves display those calculated values.

7.3 Simulated Distance and Slant Contrast

In Experiment 1, the simulated distance to the stimulus was varied from 38 to 570 cm. We included a manipulation of distance because the reliability of stereoscopic signals decreases with increasing distance; in particular, the weight given the stereo-based slant estimator ought to decrease and the weight given the nonstereo-based estimator ought to rise (Sedgwick, 1986; Backus & Banks, 1999; Johnston, et al, 1993). For those reasons, slant contrast might vary in interesting ways with simulated distance.

There was indeed a clear effect of distance on the inducer's estimated slant in the inconsistent condition of Experiment 1. With increasing distance, the inducer's perceived slant lessened relative to its stereo-specified slant (upper panels, Figures 4 and 5) and so the nonstereo weight increased systematically with distance (Tables 1 and 2). There was also a small, but systematic effect of distance on the magnitude of slant contrast *per se:* the lower panels of Figures 4 and 5 show that the slopes of the null-setting functions in the inconsistent condition increased with distance. As suggested by the discussion leading up to equation (6B), the slant contrast effect is large when the weights given to the inducer's nonstereo-specified slant (w_{ip}) and to the indirect term ($w_{t,ind}$) are both high (equation (3A)). We know that w_{ip} should increase with distance, but what about $w_{t,ind}$? The term weighted by $w_{t,ind}$ contains the relative disparity gradient ($\hat{S}_{rel}$) whose reliability should decrease with distance. However, the reliability of the competing slant estimator–the test strip's stereo-based estimator (we are concerned with $\hat{S}_{td}$ and not $\hat{S}_{tp}$ because the nonstereo estimator is unreliable due to the sparse, random texture)–should decrease in similar fashion (see discussion of Figure 11). Thus, we presume that the weight given the indirect term should change little with distance.

A stronger effect of distance should be observed if the test strip were given a regular texture. In that case, the strip's perceived slant would be affected by the nonstereo estimator. As distance increased and the weight to that estimator rose, the test strip's perceived slant should be affected less and less by the inducer and, as a consequence, slant contrast should lessen.

7.4 Does the Test Strip Affect the Perceived Slant of the Inducer?

We assumed in the development of slant estimation theory that the inducer affects the perceived slant of the test strip, but not the reverse. This assumption is supported by an observation by van Ee (1995) (see also Footnote 8). He replaced the test strip with a small, regularly textured patch superimposed (transparently) on a large, cross-hatched inducer. The patch was frontoparallel and the inducer's slant was specified stereoscopically (nonstereo slant was zero). Although the inducer was a plane, observers reported that it appeared to have a wrinkle in the region of the test patch: the inducer's apparent slant was locally distorted in a direction opposite to the patch's apparent slant. This effect can be explained by slant estimation theory. Adding a regular texture to the test patch enriches its nonstereo-specified slant and increases the weight given to the slant estimate from the patch's signals alone (the first term in equation (3A)); the patch's perceived slant is thereby reduced. Moreover, the increased reliability of the signals from the test patch increases the weight given to the indirect term in the estimate of the inducer's slant (the second term in equation (3B)); the test patch ends up having a significant effect on the inducer's perceived slant. The local nature of the perceived wrinkle suggests, not surprisingly, that the relative disparity gradient is measured locally.

We were able to replicate this effect–a perceived wrinkle in the inducer–in our setup by placing a regular texture on the test strip. We did not observe the effect when the strip's texture was random, so for the stimuli used in the experiments reported here, the test strip apparently had no discernible effect on the perceived slant of the inducer.

7.5 Stimulus Duration

Werner (1937) and Kumar and Glaser (1993) reported that the magnitude of slant contrast decreases with stimulus duration. For the stimuli they used, slant estimation theory claims that perceived slant is given by equation (6B): $\hat{S}_t = -w_{ip} S_{id}$. Thus, a decrease in slant contrast would occur if the weight w_{ip} decreased over time; the weight w_{id} would have to increase over time because we assumed that the weights add to 1. There is some evidence that the weight given stereo-specified slant increases over time relative to the weight given nonstereo-specified slant. For example, the estimated slant of a surface increases with duration when the slant is defined by disparity (and nonstereo signals indicate that it is frontoparallel; Gillam, Chambers, & Russo, 1988; van Ee & Erkelens, 1996a).

Pierce et al. (1998) claimed that little, if any, slant contrast occurs with horizontal-axis slants. Stimulus duration was unlimited in their experiment, so the failure to observe slant contrast might be due to the increasing weight given stereo-specified slant over time.

7.6 Computer Displays vs Real Surfaces

The magnitude of slant contrast was affected profoundly by the manner in which the stimuli were displayed. When we made the inducer's stereoscopic and nonstereoscopic signals compatible in images on a display screen (consistent condition, Experiment 1), we observed a reduction in slant contrast. When we presented real planes with nominally the same properties as the computer displays (Experiment 2), slant contrast was eliminated. The critical difference was presumably the presence of uncontrollable slant cues in the computer displays. Those cues include the blur gradient across the display screen and the graininess created by the display's pixelization, both of which specify the slant of the screen rather than the slant of the computed image. Apparently, the weight given the uncontrollable nonstereo signals rivals the weight given the controllable nonstereo signals (see Table 3). This result is an important reminder that surface cues from computer displays provide signals that may well affect interpretation of the visual stimulus (Banks & Backus, 1998; Stevens & Brookes, 1988; Ryan & Gillam, 1994).

8 Conclusion

In the slant-contrast effect, a small frontoparallel surface is perceived as slanted when it is surrounded by a larger slanted surface. We presented a general theory of slant estimation from which one can understand slant contrast. The theory determines surface slant via linear combination of various slant estimators; the weight of each estimator is proportional to its reliability. According to the theory, slant contrast occurs when two conditions are present: 1) the stereo- and nonstereo-specified slants of the inducing surface differ and 2) the absolute slant of the inducer and the relative slant between test strip and inducer are both estimated with greater reliability than the absolute slant of the test strip. The theory's predictions were confirmed in three experiments. We conclude that slant contrast is a byproduct of the visual system's reconciliation of conflicting information while it attempts to determine surface slant.

Acknowledgements. This work was supported by research grants from AFOSR (93NL366) and NSF (DBS-9309820). RVE was supported by Human Frontier of Science (RG-34/96) and the Foundation for Life Sciences (SLW Talent Stipendium, NWO #810-404-006/1) of the Netherlands Organization for Scientific Research. BTB was partially supported by NIH (T32 EY0704-18). We thank Payam Saisan for technical assistance, Jacquelin van Ee for long hours as an observer, and two reviewers for thoughtful suggestions.

References

Anstis, S.M. (1975). What does visual perception tell us about visual coding? In Blakemore, C., & Gazzaniga, M.S. (Eds), *Handbook of Psychobiology*, (pp. 269–323). New York: Academic Press.

Anstis, S. M., Howard, I. P. & Rogers, B. (1978). A Craik-O'Brien-Cornsweet illusion for visual depth. *Vision Research,* 18, 213-217.

Backus, B.T. & Banks, M.S. (1999). Estimator reliability and distance scaling in stereoscopic slant perception. *Perception,* in press.

Backus, B.T., Banks, M.S., van Ee, R. & Crowell, J.A. (1999). Horizontal and vertical disparity, eye position, and stereoscopic slant perception. *Vision Research,* 39, 1143–1170.

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.

Brookes, A. & Stevens, K.A. (1989). The analogy between stereo depth and brightness. *Perception,* 18, 601–614.

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.

Clarke, J.J. & Yuille, A.L. (1990). *Data Fusion for Sensory Information Processing Systems.* Boston: Kluwer.

Ebenholtz, S.M. & Paap, K.R. (1973). The constancy of object orientation: Compensation for ocular rotation. *Perception & Psychophysics,* 14, 458–470.

Erkelens, C.J. & van Ee, R. (1998). A computational model of depth perception based on headcentric disparity. *Vision Research,* 38, 2999–3018.

Fahle, M. & Westheimer, G. (1988). Local and global factors in disparity detection of rows of points. *Vision Research,* 28, 171–178.
Gillam, B. (1967). Changes in the direction of induced aniseikonic slant as a function of distance. *Vision Research,* 7, 777–783.
Gillam, B., Chambers, D. & Russo, T. (1988). Postfusional latency in stereoscopic slant perception and the primitives of stereopsis. *Journal of Experimental Psychology: Human Perception and Performance,* 14, 163-175.
Gillam, B., Flagg, T. & Finlay, D. (1984). Evidence for disparity change as the primary stimulus for stereoscopic processing. *Perception & Psychophysics,* 36, 559-564.
Gillam, B. & Lawergren, B. (1983). The induced effect, vertical disparity, and stereoscopic theory. *Perception & Psychophysics,* 34, 121-130.
Gillam, B. & Ryan, C. (1992). Perspective, orientation disparity, and anisotropy in stereoscopic slant perception. *Perception,* 21, 427-439.
Gogel, W.C. (1956). The tendency to see objects as equidistant and its inverse relation to lateral separation. *Psychological Monographs,* 70 (Whole No. 411), 1–17.
Graham, M.E. & Rogers, B. J. (1982). Simultaneous and successive contrast effects in the perception of depth from motion-parallax and stereoscopic information. *Perception,* 11, 247–262.
Harker, G.S. (1962). Apparent frontoparallel plane, stereoscopic correspondence, and induced cyclorotation of the eyes. *Perceptual and Motor Skills,* 14, 75–87.
Howard, I.P. & Kaneko, H. (1994). Relative shear disparities and the perception of surface inclination. *Vision Research,* 34, 2505-2517.
Howard, I.P., Ohmi, M., & Sun, L. (1993). Cyclovergence: A comparison of objective and psychophysical measurements. *Experimental Brain Research,* 97, 349-355.
Howard, I.P. & Rogers, B.J. (1995). *Binocular vision and stereopsis.* New York: Oxford.
Johnston, E.B., Cumming, B. G. & Parker, A. J. (1993). Integration of depth modules: Stereopsis and texture. *Vision Research,* 33, 813-826.
Kumar, T. & Glaser, D.A. (1991). Influence of remote objects on local depth perception. *Vision Research,* 31, 1687-1699.
Kumar, T. & Glaser, D.A. (1992). Shape analysis and stereopsis for human depth perception. *Vision Research,* 32, 499-512.
Kumar, T. & Glaser, D.A. (1993). Temporal aspects of depth contrast. *Vision Research,* 33, 947-957.
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. & Longuet-Higgins, H. C. (1982). A computational model of binocular depth perception. *Nature,* 297, 376-378.
Mitchison, G.J. & McKee, S.P. (1990). Mechanisms underlying the anisotropy of stereoscopic tilt perception. *Vision Research,* 30, 1781-1791.
Mitchison, G.J. & Westheimer, G. (1984). The perception of depth in simple figures. *Vision Research,* 24, 1063-1073.
Mitchison, G.J. & Westheimer, G. (1990). Viewing geometry and gradients of horizontal disparity. In Blakemore, C., (Eds), *Vision: Coding and Efficiency*, (pp. 302-309). Cambridge: Cambridge University Press.
Nelson, J. (1977). The plasticity of correspondence: after-effects, illusions and horopter shifts in depth perception. *Journal of Theoretical Biology,* 66, 203-266.
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. (1946). The binocular depth contrast phenomenon. *American Journal of Psychology,* 59, 111-126.
Pastore, N. (1964). Induction of stereoscopic depth effect. *Science,* 144, 888.
Pastore, N. & Terwilliger, M. (1966). Induction of stereoscopic depth effects. *British Journal of Psychology,* 57, 201-202.
Pierce, B.J., Howard, I.P. & Feresin, C. (1998). Depth interactions between inclined and slanted surfaces in vertical and horizontal orientations. *Perception,* 27, 87-103.
Porrill, J., Mayhew, J.E.W. & Frisby, J.P. (1990). Cyclotorsion, conformal invariance, and induced effects in stereoscopic vision. In Ulman, S. & Richards, W, (Eds), *Image Understanding.* Norwood: Ablex.
Richards, W. (1972). Response functions for sine- and square-wave modulations of disparity. *Journal of the Optical Society of America,* 62, 907-911.

Rogers, B.J. & Bradshaw, M. (1995). Disparity scaling and the perception of frontoparallel surfaces. *Perception,* 24, 155-179.
Rogers, B.J. & Graham, M. E. (1983). Anistropies in the perception of three-dimensional surfaces. *Science,* 221, 1409-1411.
Ryan, C. & Gillam, B. (1994). Cue conflict and stereoscopic surface slant about horizontal and vertical axes. *Perception*, 23, 645-658.
Sato, M. & Howard, I.P. (1999). Effects of perspective on stereoscopic depth contrast of inclined surfaces. *Submitted.*
Sedgwick, H.A. (1986). Space perception. In Boff, K.R., Kaufman, L. & Thomas, J.P., (Eds), *Handbook of Perception and Human Performance*, (Ch. 21 pp.1-57). New York: Wiley.
Schumer, R.A. & Ganz, L. (1979). Independent stereoscopic channels for different extents of spatial pooling. *Vision Research,* 19, 1303-1314.
Stevens, K.A. (1983). Slant-tilt: the visual encoding of surface orientation. *Biological Cybernetics,* 46, 183-195.
Stevens, K.A. & Brookes, A. (1988). Integrating stereopsis with monocular interpretations of planar surfaces. *Vision Research,* 28, 371-386.
Stevenson, S.B., Cormack, L.K. & Schor, C.M. (1991). Depth attraction and repulsion in random dot stereograms. *Vision Research,* 31, 805-813.
van Ee, R. (1995). *Stability of binocular depth perception*. Utrecht University. ISBN 90-393-0846-2.
van Ee, R. & Erkelens, C.J. (1996a). Temporal aspects of binocular slant perception. *Vision Research,* 36, 43-51.
van Ee, R. & Erkelens, C.J. (1996b). Anisotropy in Werner's binocular depth contrast effect. *Vision Research,* 36, 2253-2262.
van Ee, R. & Erkelens, C.J. (1996c). Stability of binocular depth perception with moving head and eyes. *Vision Research,* 36, 3827-3842.
Werner, H. (1937). Dynamical theory of depth perception. *Psychological Monographs,* 49, 1-127.
Werner, H. (1938). Binocular depth-contrast and the conditions of the binocular field. *American Journal of Psychology,* 51, 489-497.
Westheimer, G. & Levi, D.M. (1987). Depth attraction and repulsion of disparate foveal stimuli. *Vision Research,* 27, 1361-1368.