# Stereoscopic 3-D content appears relatively veridical when viewed from an oblique angle

**Paul Hands** Institute of Neuroscience, Newcastle University, Tyne and Wear, UK

**Tom V. Smulders** Institute of Neuroscience, Newcastle University, Tyne and Wear, UK

**Jenny C. A. Read** Institute of Neuroscience, Newcastle University, Tyne and Wear, UK

**Geometrically, stereoscopic 3-D (S3D) content should appear distorted unless viewed from the position for which the content was produced. Almost all commercial and laboratory S3D content is generated assuming that it will be presented on a screen frontoparallel to the viewer. However, in cinema and the home, S3D content is regularly viewed from oblique angles, and yet shapes are not usually perceived to be distorted. It is not yet known whether this is simply because viewers are insensitive to incorrect viewing angles or because viewers automatically compensate for oblique viewing, as they do for 2-D content. Here, we investigate this using a canonical-form paradigm. We show that S3D content can indeed appear warped when viewed from oblique angles, and that this effect is more pronounced than for 2-D content. We hypothesized that motion cues in the content would aid in the correct perception of S3D content, making it appear more natural even when viewed obliquely, but we find little support for this idea. However, the perceptual distortions are still small, and viewers do compensate to some extent for oblique viewing. We conclude that, at least as regards object distortion, oblique viewing is unlikely to be substantially more of a problem for S3D content than it already is for 2-D.**

## Introduction

Due to the horizontal offset between the eyes, they receive slightly different retinal images. These small binocular disparities are detected by the brain and, even in the absence of other depth cues, suffice to create a vivid perception of depth (Wheatstone, 1838). This effect is exploited in stereoscopic displays, which present separate images to the two eyes. Cinema, home television systems with stereoscopic 3-D (S3D) capabilities, and some game consoles use different types of S3D displays, including passive and active stereo and parallax barrier technology.

S3D displays make it possible to re-create the different retinal images caused by a real object in space. This exact recreation is often referred to as orthostereoscopic, or orthostereo (Kurtz, 1937). However, almost no commercial S3D content is orthostereoscopic. To display in S3D orthostereoscopically, it is necessary to control and coordinate all the aspects of the content production, from capture to display. Mathematically, S3D displays produce an orthostereo image only when the viewer is positioned with each eye exactly at the center of projection for which that eye's image was filmed or rendered (Held & Banks, 2008; Woods, 1993). If the viewer moves away from this specified position, the object depicted by the retinal stimulus will alter. Indeed, given the position of the eyes, the retinal disparities will in general be non-epipolar, i.e., not consistent with any physical object (Held & Banks, 2008; Read, Phillipson, & Glennerster, 2009; Woods, 1993).

We can distinguish two main ways in which viewers can move away from the center of projection. First, they may view content from the wrong distance. Second, they may view content from the wrong angle. Previous studies have shown that incorrect viewing distance can lead to distortions in perceived depth and shape (Held & Banks, 2008; Woods, 1993). Sometimes there is no correct viewing distance. Commercial S3D content is often filmed with the cameras "toe-in," i.e., converged on the object of interest. This produces a "keystone" distortion in the images. To be orthostereo,

Citation: Hands, P., Smulders, T. V., & Read, J. C. A. (2015). Stereoscopic 3-D content appears relatively veridical when viewed from an oblique angle. *Journal of Vision, 15*(5):6, 1–21, http://www.journalofvision.org/content/15/5/6, doi:10.1167/15.5.6.

doi: 10.1167/15.5.6 Received May 14, 2014; published April 17, 2015 ISSN 1534-7362 © 2015 ARVO

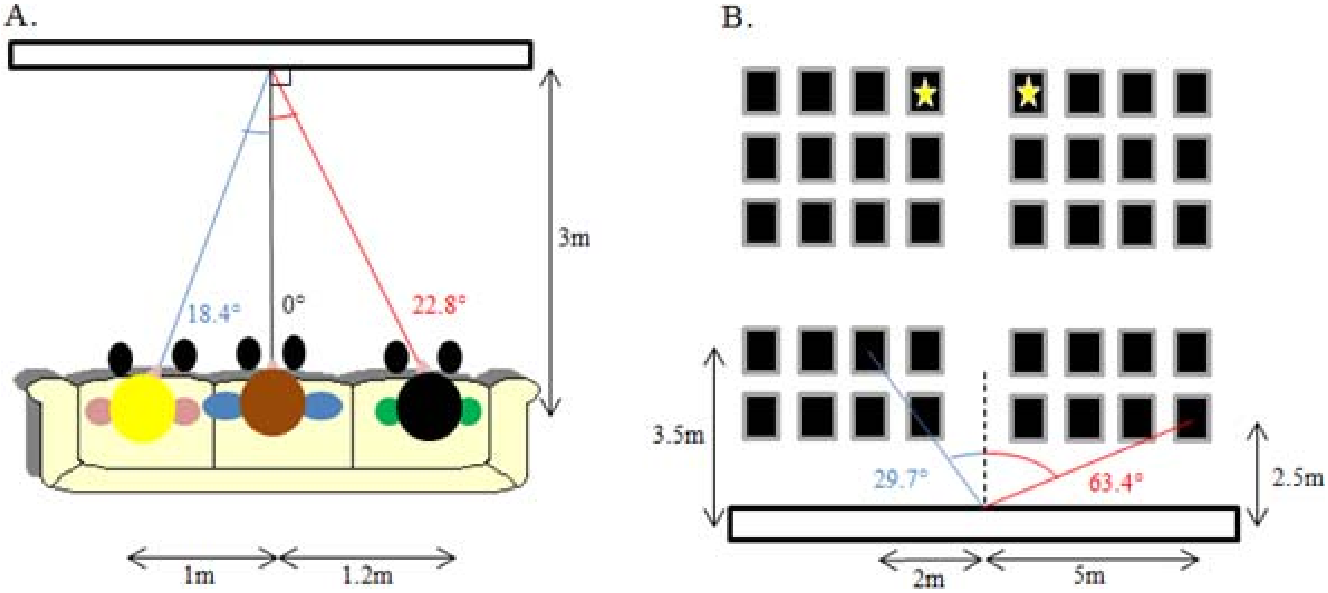


Figure 1. (A) A diagram to illustrate how viewing angles can change. If three people sit in the home watching television together, 3 m away from the screen and between 1 and 1.2 m apart, the viewing angle can be up to 23°. Angles are measured from perpendicular viewing (here we show a top-down diagram). (B) Viewing angles for different seats in a hypothetical cinema theater auditorium. Stars indicate the seats with viewing angles closest to ideal (0° = frontoparallel).

such content has to be either corrected for the distortion or viewed on two screens: one for each eye, orthogonal to the line of sight from the respective eye. While this can be arranged in a laboratory haploscope, it is almost never the case for commercial S3D. If uncorrected content filmed with converged cameras is presented on a single screen, the pattern of vertical disparities could only occur in reality if the viewers' eyes were more converged than is the case when they view the content (Banks, Read, Allison, & Watt, 2012). Thus, there is no viewing position for which the content is orthostereo. In any case, the correct viewing distance will typically vary during a feature. In a mass-viewing venue like a cinema, viewing distance will vary greatly for different audience members.

Viewing angle is more straightforward, in that there is a clear "correct" viewing angle: Almost all S3D content is created to be viewed on a screen frontoparallel to the viewer. More specifically, the eyes should be positioned such that the plane bisecting the interocular axis is normal to the screen and passes through the center of the screen. However, both in cinemas and at home, many viewers will be viewing the screen obliquely. Even if they turn their head towards the center of the screen, such that the plane bisecting the interocular axis passes through the screen midline, this plane will not be normal to the screen (Figure 1A). Similarly, in a cinema theater, viewers seated at the extreme front and side of the auditorium will be subject to a very large deviation away from the perpendicular viewpoint (Figure 1B). This is problematic since, geometrically, the shape specified by a 3-D display changes with the viewing angle (Held & Banks, 2008; Woods, Docherty, & Koch, 1993). Thus, if human depth perception were based on the geometry of the retinal images, content created to be viewed perpendicularly should look distorted from any other viewing angle.

Of course, these problems also apply to 2-D images, in the sense that the image projected onto the retina varies as a function of viewing angle. The problem of why images nevertheless appear veridical from a range of viewing angles has fascinated researchers since the Renaissance (Kubovy, 1988; Pirenne, 1970). Several factors seem to contribute. One is that humans are simply not very sensitive to the distortion introduced by oblique viewing (Cutting, 1987; Gombrich, 1972). Additionally, images usually depict familiar objects, so that viewers' perceptions can be influenced by their expectations (Thouless, 1931). However, it is also clear that observers are capable of compensating for the oblique viewing, so that perception is based not on the image actually projected onto the retina but on the image which would have been seen if viewed perpendicularly (Perkins, 1973; Rosinski, Mulholland, Degelman, & Farber, 1980; Vishwanath, Girshick, & Banks, 2005). This compensation could work by recovering the true center of projection and reinterpreting the retinal image accordingly. The true center of projection could be estimated from cues present within the depicted scene (De La Gournerie, 1859; Saunders & Backus, 2007), such as the location of vanishing points, or from external cues regarding the orientation of the picture plane combined with simplifying assumptions such that the true center of projection lies on the central surface normal. Presumably, such a mechanism would have to reflect experience with 2-D pictures (Deregowski, 1969; Jahoda & McGurk, 1974a, 1974b; Olson & Boswell, 1976). Vishwanath et al. (2005) have recently argued for a simpler heuristic, whereby the retinal image is reinterpreted locally based on local surface slant. They argue that this may reflect a more general heuristic which is useful when interacting with real objects viewed obliquely, not a specific mechanism for interpreting pictures. External cues to local surface slant include binocular disparity, vergence, accommodation, the position of specular highlights relative to external light sources, and perspective cues provided by a frame surrounding the screen plane. Accordingly, occluding the frame of the display, viewing monocularly, or viewing through a pinhole all tend to make the compensation less effective, so that images appear warped when viewed at oblique angles (Bereby-Meyer, Leiser, & Meyer, 1999; Perkins, 1973; Vishwanath et al., 2005).

There is a widespread belief that this compensation process is less effective for S3D stimuli (Banks, Held, & Girshick, 2009; Bereby-Meyer et al., 1999; Perkins, 1973; Pirenne, 1970; Zorin & Barr, 1995). There are several reasons why this should be so. In 2-D displays, disparity and vergence are powerful cues which specify that the picture lies on a flat plane and also indicate the orientation of this plane. Critically, these binocular cues are unaffected by the contents of the picture and therefore allow the viewer to estimate screen slant without confounds. In S3D, both these cues now indicate that the scene is not planar but consists of objects at different depths (Bereby-Meyer et al., 1999). In the words of Pirenne (1970, p. xiv+199), "in the case of [stereoscopic images], the observer is hardly aware of the surface of the picture, as a surface." Ironically, therefore, the very thing that makes S3D a powerful visual experience, namely the use of binocular disparity to depict 3-D objects in space rather than lying on a flat picture plane, might make viewers less able to correct for oblique viewing. Additionally, despite recurrent upsurges of interest in S3D displays since the 19th century, viewers will have had far less exposure to S3D pictures than to 2-D. If experience with 2-D pictures plays a role in compensating for oblique viewing, these mechanisms may not have developed to the same extent for S3D.

Surprisingly, however, this widespread belief has been little tested. We are aware of only three previous studies other than our own which have considered perceptual distortions in stereoscopic 3-D due to oblique viewing (Banks et al., 2009; Bereby-Meyer et al., 1999; Perkins, 1973). The study by Banks et al. (2009) is the only one to compare perception of 2-D and S3D stimuli, although only one observer viewed both. The researchers concluded that, as predicted, percepts from stereo pictures are significantly more affected by oblique viewing angle than are percepts from conventional, 2-D pictures.

All three previous studies used static content. This is a potentially serious limitation, given that commercial S3D usually consists of video content, which contains powerful internal structure-from-motion cues. There are good theoretical reasons for expecting that these cues could affect viewers' ability to compensate for oblique viewing angle (Cutting, 1987). The interpretations of 3-D shapes based on motion are underdetermined: The sequence of images is consistent with many possible movements of objects in the world. Thus, humans need to apply additional constraints, such as the rigidity assumption: "Any set of elements undergoing a 2D transformation which has a unique interpretation as a rigid body moving in space, should be interpreted as such" (Ullman, 1979, p. 411). Humans are very good at reconstructing this interpretation when they view a series of such 2-D images. However, when the same series of frames is viewed obliquely, the successive retinal images will not in general be geometrically consistent with a rigid body in motion. Mathematically, this is the same phenomenon discussed before for stereographic 3-D: A stereogram designed to be orthostereographic for frontoparallel viewing becomes nonepipolar—geometrically inconsistent with any object—when viewed obliquely (Held & Banks, 2008). In stereo, the visual system is capable of extracting the nonepipolar component of disparity and using it to change the interpretation of the epipolar component, effectively interpreting the scene as if it were being viewed with a different eye position (Mayhew & Longuet-Higgins, 1982; Ogle, 1938; Rogers & Bradshaw, 1993). Conceivably, a related computation might be present in the motion domain: The visual system might be able to use the rigidity assumption to estimate the angle from which a projected image is being viewed, as well as the shape of the object and its motion relative to the eye. As we have seen, in picture perception the brain has to decide whether it is viewing a projection of Shape 1 from the correct angle or a projection of Shape 2 from an incorrect viewing angle. We have already seen some ways the visual system might in principle choose between these, e.g., by using disparity cues from the picture surface to deduce that the viewing angle is incorrect. However, with a dynamic stimulus the brain has to decide whether it is viewing a projection of a moving, deforming Shape 1 from the correct angle or a projection of a moving, rigid Shape 2 from an incorrect angle. An assumption that objects are generally rigid would tend to result in the latter choice. Since the rigidity assumption would apply equally to 2-D and S3D content, this would tend to reduce the difference between S3D and 2-D content

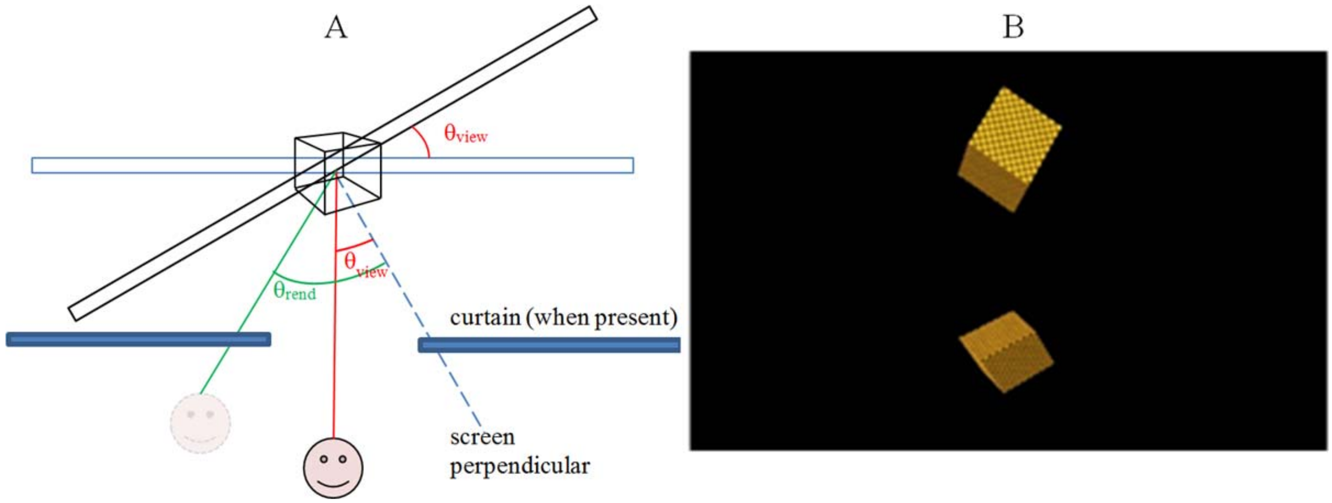


Figure 2. Experimental setup seen from above (A) and example stimulus (B). (A) The heads of two possible viewers are sketched. The head drawn with solid lines represents the actual position of the participant, whose line of sight to the screen (red line) is at angle $\theta_{\text{view}}$ to the perpendicular (blue dashed line). The head drawn with dotted lines represents the position of a hypothetical viewer whose line of sight is at $\theta_{\text{rend}}$ (green line). The obliquely rendered cube is projected correctly for this hypothetical viewer. It is correct for the participant only when $\theta_{\text{view}} = \theta_{\text{rend}}$. The normal-rendered cube is projected correctly for a second hypothetical viewer whose line of sight is perpendicular to the screen. In some blocks, a curtain was pulled across so as to occlude the edges of the screen from the participant's view. (B) Stimulus drawn on the display screen. Here, the top cube is the normal-rendered cube (rendered for viewing perpendicular to the screen) and the bottom cube is the obliquely rendered cube, here rendered for a viewing angle of 45°. You may find that the bottom cube appears less distorted when viewed from 45° to the right.

otherwise expected from the disrupted binocular cues in S3D.

In the present study, we addressed this question using a canonical-form task in which subjects were asked to report their perception of cubes rendered for perpendicular and oblique viewing. Cubes are a familiar object which has been used in many previous studies of picture perception (Cutting, 1987; Hagen & Elliott, 1976; Hagen, Elliott, & Jones, 1978; Perkins, 1973). In a previous study (Hands & Read, 2013), we used static wire-frame cubes. These displayed the well-known Necker illusion (Necker, 1832), i.e., they could be perceived in one of two different orientations. Because our cubes were rendered using perspective projection and presented fairly close to the observer, only one of the two orientations appeared as a cube; the other appeared as a warped frustum. To avoid distortions caused by this effect, in the present study we used solid cubes (Figure 2B) whose orientation was unambiguous. The cubes were covered with a checkerboard pattern. The stimuli thus contained several cues which could potentially be used to judge whether the objects were perfect cubes with parallel equal-length sides and right-angle corners (e.g., perspective, shading, texture). We examined the effect of three factors on perceptual compensation. To determine the effect of the picture frame, we compared results when the edges of the screen were occluded versus when they were visible. To examine motion, we interleaved static objects with objects depicted as rigidly rotating. Using an S3D display, we interleaved monocular, binocular 2-D, and stereoscopic 3-D stimuli to test whether the visual system is less able to compensate for oblique viewing in S3D than in 2-D.

## Material and methods

### Participants

Participants were recruited via an internal volunteer scheme at Newcastle University's Institute of Neuroscience, on the basis that they had no visual problems other than wearing glasses or contact lenses. The work was approved by the Newcastle University Faculty of Medical Sciences Ethics Committee. Ten participants (nine women, eight of them naïve to the purposes of the study; one man, PH) were used in the study. Only one voluntary participant and the author took part in both the previous study (Hands & Read, 2013) and the current one, due to availability of the other participants. Naïve participants were not informed of the experimental aims or hypotheses, but due to the random order of blocks they will have been able to work out that the viewing angle was changing. Participants were paid £10 for completing the study.

## Apparatus

Stimuli were presented on a 50-in. stereoscopic 3-D monitor (LG 47LD920-ZA) using passive stereo technology. The resolution of the monitor was 1920 pixels wide × 1080 high, and left/right eye images are presented on alternate pixel rows, so that each image has a vertical resolution of 540 pixels. As described later, the monitor was used in 2-D mode to avoid artifacts due to the vertical averaging performed by the monitor in 3-D mode. The maximum luminance of the display was 20 cd/m$^2$, as measured through the 3-D glasses with a Minolta LS100 photometer. Interocular cross talk was 1.4% when measured with the screen frontoparallel to the photometer, rising to 2.0% for a viewing angle of 20° and 7.1% for a viewing angle of 45°.

Participants sat at a viewing distance of 120 cm, measured perpendicularly from the center of the screen to the midpoint of the eyes, with their eyes at the same height as the center of the screen. They wore passive 3-D glasses throughout the experiment, enabling us to interleave S3D, 2-D, and monocular stimuli. The monitor sat on a turntable, which allowed it to be accurately rotated between ±45° about a vertical axis passing through the midline of the screen. We define the viewing angle $\theta_{\text{view}}$ to be the angle between the plane normal to the screen and the viewer's line of sight to the center of the screen (Figure 2A). In different experimental blocks, the turntable was rotated so that $\theta_{\text{view}}$ was either 0°, –45° (closer to the viewer on her right), or +20°. It was convenient to alter the viewing angle by moving the display screen rather than the participant (see Figure 2A). A chin rest was used to ensure that the participant's eyes were at the correct position, and the chair was adjustable to ensure that the participant was comfortable. In some experimental blocks, a fabric curtain with a hole was pulled across that occluded all four screen edges from the participant's view while allowing them to see the stimuli.

## Stimulus generation

Stimuli were generated and the experiments run using the computer programming environment MATLAB (MathWorks, Natick, MA) and the Psychtoolbox extension (Brainard, 1997; Kleiner, Brainard, & Pelli, 2007; Pelli, 1997). For each frame of the stimulus, we generated separate left and right images of resolution 1920 × 540 pixels, treating each pixel as being effectively a rectangle twice as high as broad (e.g., a frame 100 pixels wide by 50 pixels high would appear square on the screen). We used the interleaved line stereomode of Psychtoolbox to combine these images on alternate pixel rows and displayed the result as a single image with the monitor in 2-D mode.

In all our experiments, virtual cubes were rendered onto the screen via central perspective projection. The center of each cube lay in the screen plane. Usually when one renders a scene, the projection plane is perpendicular to the line from the center of projection to the center of the scene. In our experiments, the projection plane was sometimes rotated away from this position (Figure 2A). To find where to render a point on this rotated projection plane, we imagine drawing a straight line from the center of projection through the point in question. The point where this line intersects the projection plane is where the point should be rendered. For a monocular viewer whose eye is a pinhole at the center of projection, this should produce exactly the same retinal image as the real object.

In the previous study (Hands & Read, 2013) using wire-frame cubes, we wrote our own MATLAB software to calculate where to render the vertices of each cube. We checked our calculations by drawing a square onto a sheet of acetate and mounted it on a sheet of Perspex in front of the screen, representing one face of the virtual cube. We supplied our code with the physical coordinates of this square and rendered it for different viewing angles. We verified that, in each case, the image drawn on the screen lined up with the physical square drawn on the acetate, confirming that our code was rendering the virtual code correctly whether the screen was perpendicular to the viewer or viewed obliquely at the specified angle. In the experiments reported here, we used Psychtoolbox with the OpenGL library to draw solid, textured cubes. We confirmed that this produced the same vertex positions by using Psychtoolbox to draw dots on top of the rendered cubes at the locations of the vertices as calculated by our own code and checking that these dots lay on the vertices of the rendered cubes.

Figure 3A shows the same wire-frame cube rendered for render angles of 0° (red) and 45° (blue). In the S3D condition, stimuli were rendered separately for left and right eyes. In the M2D (monocular) condition, one eye saw the same stimulus as in the S3D condition while the other eye saw a black screen (except for any cross talk). In the B2D (binocular) condition, the stimulus was rendered as if for a single cyclopean eye in the middle of the two actual eyes. We used a standard interocular-distance value of 6.3 cm, close to the average for adult humans (Dodgson, 2004). Commercial S3D content is necessarily generated for a standard viewer, and we were interested in measuring the effect of oblique viewing under these conditions. Additionally, our data indicate that viewers are insensitive to large errors in the angle at which they view the screen (>10°), so it seems unlikely they were very sensitive to errors caused by the small variation in interocular distance.

A)

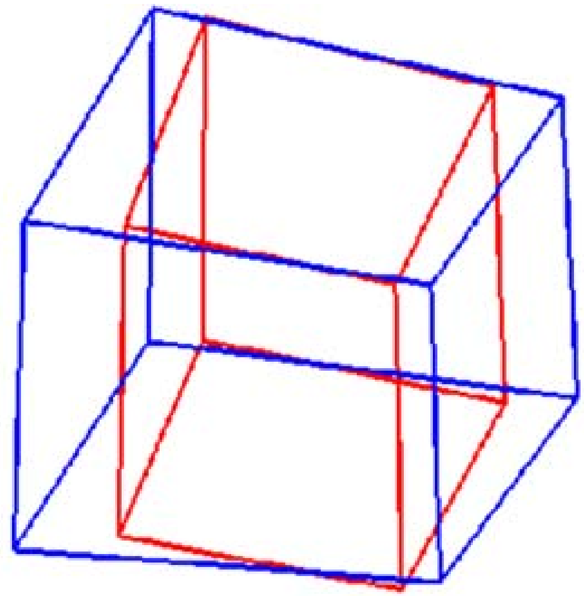

B)

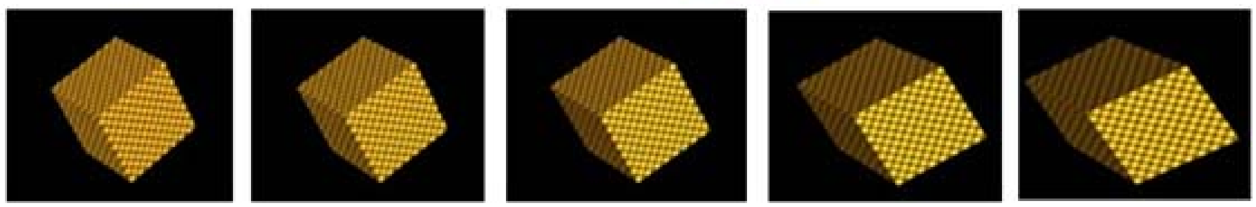

Figure 3. (A) Rendered image on-screen for two different render angles. The two cubes are projections from exactly the same virtual cube in space, but the red cube is rendered for perpendicular viewing, while the blue cube is rendered for a viewing angle of 45°. If you view the image with one eye from 45° to the right, the blue cube should appear as the red cube does when viewed normally. (B) Example cube rendered for five different angles used in our experiments. From left to right, $\theta_{\text{rend}} = 0°$, 10°, 20°, 35°, and 45°.

## Experimental design

In each trial, the participants viewed two cube-like objects, one rendered onto the top half of the screen and one onto the bottom. The participants were asked to choose which cube looked the "most cube-like" in the sense of having equal-length sides and all right-angle vertices. They indicated their answer by pressing the up or down arrow on the keyboard.

The objects were perspective projections of virtual cubes in space. The center of the virtual cube was always in the screen plane, one quarter screen height either above or below the center of the screen. This was unaffected by the screen orientation, since the midline of the screen was the axis of rotation.

In each trial, one of the two cubes was rendered for frontoparallel viewing in the normal way, i.e., for a line of sight perpendicular to the screen. The other was rendered for an oblique viewing angle that varied between $\theta_{\text{rend}} = -45°$ and $\theta_{\text{rend}} = +45°$. We will refer to these as the normal-rendered and obliquely rendered cube, respectively. When $\theta_{\text{rend}} = \theta_{\text{view}}$, the obliquely rendered cube was rendered for the actual viewing angle of the participant. We will refer to this as geometrically correct. In the S3D condition, the geometrically correct stimulus is orthostereoscopic, i.e., each eye ideally saw the retinal image which would have been projected by a physical cube in front of the viewer, apart from accommodation effects. On each trial, the orientation of each cube was random: Each virtual cube was rotated through a random angle about all three axes in succession before being rendered.

Figure 3B shows a cube rendered for five different values of $\theta_{\text{rend}}$.. The apparent distortion increases monotonically as the rendering angle departs from frontoparallel. Additionally, a given cube has a wider horizontal extent on the screen when rendered for oblique viewing (Figure 3). To help ensure that participants did not simply judge the "more cube-like" object to be the one with the smallest extent on-screen, the size of the virtual cubes was chosen randomly on each trial. The side length $L$ of one cube was picked from a uniform distribution between 6 and 14 cm, and the side length of the other cube was then set to $20 - L$ cm. The sum of the two side lengths was therefore always 20 cm, ensuring that the rendered cubes never overlapped on the screen. This manipulation meant that the obliquely rendered cube could be either larger or smaller than the normal-rendered one.

In static trials, both objects remained stationary on the screen; in motion trials, both objects rotated at a constant speed of 18°/s about all three axes (see Supplementary Movie S1). This rotation speed was chosen as being slow enough to be comfortable for the participant to follow yet fast enough to produce rapid changes in the on-screen image and thus powerful structure-from-motion cues. In both types of trials, the objects remained on-screen until the participant indicated whether the top or bottom object appeared most cube-like.

## Experimental parameters

The experiment was composed of six blocks. In each block, the participant sat at one of three viewing angles, $\theta_{\text{view}} = -45°$, 20°, or 0°, and had the curtain occluder either present or absent. In blocks where the occluder was present, it was always pulled across before the television's orientation was changed, so the participant had no prior knowledge of the screen orientation. Each participant did the six blocks in a random order chosen with a random number generator. In each block, the following four parameters were manipulated:

- The angle $\theta_{\text{rend}}$ used to project the obliquely rendered cube (eight possible values: ±45°, ±35°, ±20°, and ±10°; see Figure 3B)
- Whether the normal-rendered cube was at top or bottom of the screen (two possible values)
- Object motion (two possible values: static or rotating)
- Binocularity (four possible values: S3D [binocular; each eye sees a different image], B2D [binocular; each eye sees the same image on the screen], or M2D [monocular] left or right).

For each combination of the first three parameters, the S3D and 2-D conditions were presented four times in each block, while the monocular-left and monocular-right trials were presented twice. Thus, each block contained $8 \times 2 \times 2 \times (4+4+2+2) = 384$ trials, in a random order chosen by the computer. No difference in results was apparent between the monocular-left and monocular-right trials, so these were pooled for analysis along with cube location (top or bottom of the screen). Thus, each block effectively contained eight repetitions of each of 48 combinations of experimental parameters (8 $\theta_{\text{rend}}$ × 3 binocularity [S3D/B2D/M2D] × 2 object motion [static/rotating]). Altering the viewing and rendering angles enables us to assess the effectiveness of perceptual compensation for oblique viewing. Binocularity, object motion, and frame occlusion are the three viewing factors whose effect on compensation we wish to assess.

## Modeling

To explain our data, we developed a mathematical model which assumes that object appearance is influenced by two competing mechanisms. First, we postulated that objects appear more veridical (in this case, cube-like) when the image on the retina is consistent with a perspective projection of a real cube (geometrically correct). In our experiments, this is the case where $\theta_{\text{rend}} = \theta_{\text{view}}$. However, both our data and the existing literature indicate a second mechanism: Objects also appear more veridical when rendered for frontoparallel viewing, $\theta_{\text{rend}} = 0°$, even if the screen is in fact viewed obliquely. We assume that the perceived veridicality due to each mechanism declines according to a Gaussian function as the value of $\theta_{\text{rend}}$ moves away from the optimum, and we further assume that the perceived veridicality of the object is simply the sum of contributions from each factor. Accordingly, we model the perceived veridicality $V$ of each object as

$$V = A\exp\left(\frac{\theta_{\text{rend}}^2}{2s^2}\right) + B\exp\left(-\frac{[\theta_{\text{rend}} - \theta_{\text{view}}]^2}{2r^2}\right), \quad (1)$$

where the free parameters $s$ and $r$ determine each factor's sensitivity to $\theta_{\text{rend}}$, and $A$ and $B$ determine the relative weight of each factor. $A$ is the weight given to normal rendering, and $B$ the weight given to geometrical correctness. In our experiments, one of the cubes was always rendered for perpendicular viewing, $\theta_{\text{rend}} = 0°$. The difference in perceived veridicality between this normal-rendered cube and the obliquely rendered cube is therefore

$$\Delta V = A - A\exp\left(-\frac{\theta_{\text{rend}}^2}{2s^2}\right) + B\exp\left(-\frac{\theta_{\text{view}}^2}{2r^2}\right) - B\exp\left(-\frac{[\theta_{\text{rend}} - \theta_{\text{view}}]^2}{2r^2}\right). \quad (2)$$

When this difference is positive, the viewer perceives the normal-rendered object as most cube-like. To account for the graded chance in performance as a function of $\theta_{\text{rend}}$ and $\theta_{\text{view}}$, as well as trial-to-trial variation, we make the usual assumption that this signal is subject to internal noise, which we model as Gaussian. Without loss of generality, we set the standard deviation of the noise to 1, since this degree of freedom is already accounted for by the weights $A$ and $B$. We assume that the viewer selects the normal-rendered object as most resembling a cube whenever their noisy internal estimate of $\Delta V$ is greater than zero.

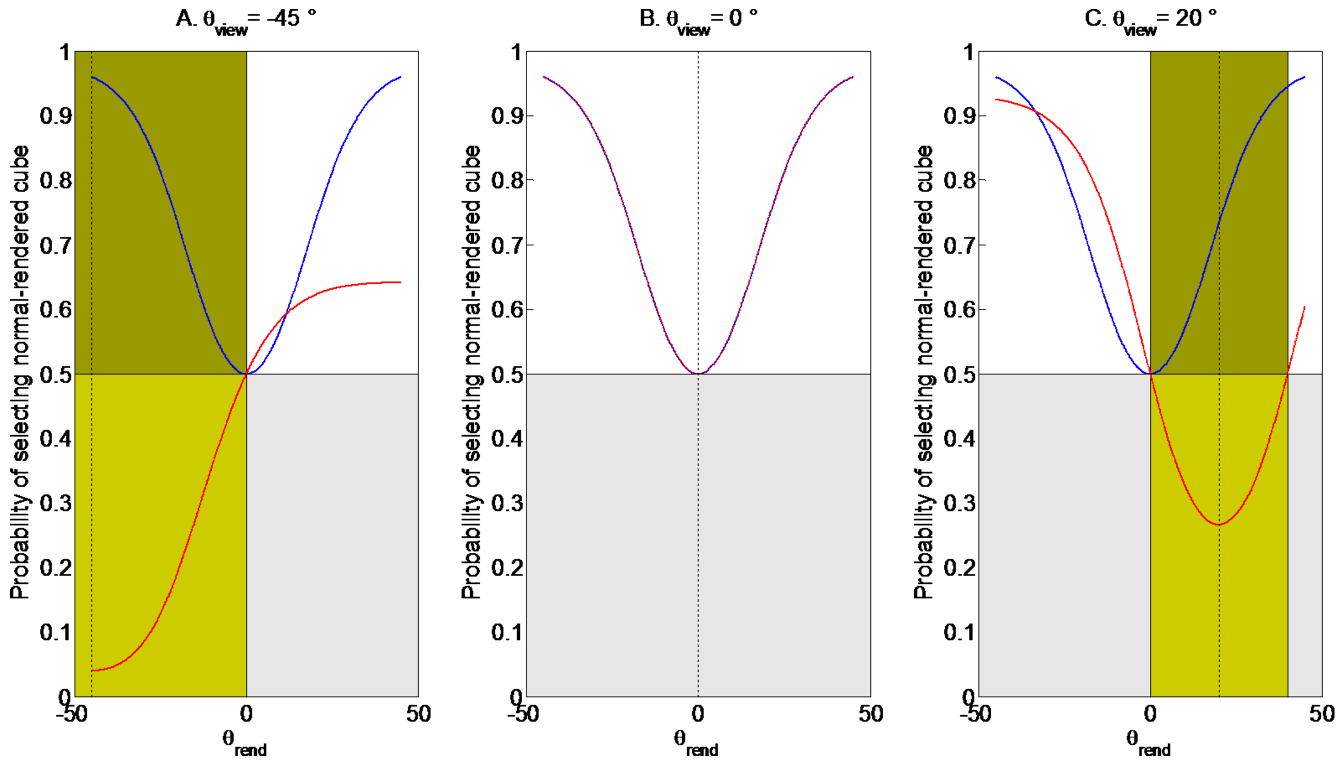


Figure 4. Model predictions with perfect compensation for oblique viewing (blue) and no compensation (red). Curves show the probability that the model selects the normal-rendered cube as being more veridical, plotted as a function of render angle $\theta_{\text{rend}}$ for three different viewing angles $\theta_{\text{view}}$. Model parameters were $r = s = 24°$; for the blue curves, $A = 3$ and $B = 0$ (perfect compensation for oblique viewing); for the red curves, $A = 0$ and $B = 3$ (no compensation). White and gray regions show where the normal-rendered cube is also closer to geometrically correct than the oblique cube ($|\theta_{\text{rend}} - \theta_{\text{view}}| > |\theta_{\text{view}}|$); yellow-shaded regions show where the obliquely rendered cube is closer to geometrically correct ($|\theta_{\text{rend}} - \theta_{\text{view}}| < |\theta_{\text{view}}|$). Light shading (white or light yellow) is used to show the direction of preference expected under the no-compensation model, i.e., below 0.5 where $|\theta_{\text{rend}} - \theta_{\text{view}}| < |\theta_{\text{view}}|$ and above 0.5 otherwise.

The probability that the viewer will select the normal-rendered object as most resembling a cube is then given by

$$P = 0.5 + 0.5 \times \operatorname{erf}\left(\frac{\Delta V}{2}\right). \tag{3}$$

At $\theta_{\text{rend}} = \theta_{\text{view}} = 0°$, the model returns a probability of 0.5 for selecting either cube, which is correct, since at this point both cubes are rendered for the same viewing angle (they would not be identical on the screen, due to the randomization of size and orientation described earlier).

To illustrate the effect of the two mechanisms, Figure 4 shows model results for two different extreme cases: perfect compensation (blue, $B = 0$) and no compensation (red, $A = 0$). With perfect compensation, the results are unaffected by viewing angle: The model always selects the normal-rendered cube when the obliquely rendered cube is rendered with a perceptibly different rendering angle. With no compensation, the model selects the obliquely rendered cube when this is closer to geometrically correct.

## Fitting

Our model assumes that the four model parameters $A$, $B$, $r$, and $s$ do not change with the viewing angle $\theta_{\text{view}}$. However, we allowed the model parameters to vary for the different viewing factors, i.e., frame occlusion, binocularity, and object motion, to account for the effect they may have on perceptual compensation. We used maximum likelihood fitting assuming simple binomial statistics, as follows. Suppose that on the $j$th set of stimulus parameters, our subjects chose the normal-rendered object on $M_j$ out of $N_j$ trials. Then the log likelihood of the data set, apart from a constant which has no effect on the fitting, is

$$\log L = \sum_j \left\{ M_j P_j + (N_j - M_j)(1 - P_j) \right\}, \tag{4}$$

where $P_j$ is the model probability for the $j$th data point, which in turn depends on the stimulus parameters $\theta_{\text{view}}$, $\theta_{\text{rend}}$, and the four model parameters, as described by Equations 2 and 3. We adjusted the model parameters to maximize this likelihood. The mathematical properties

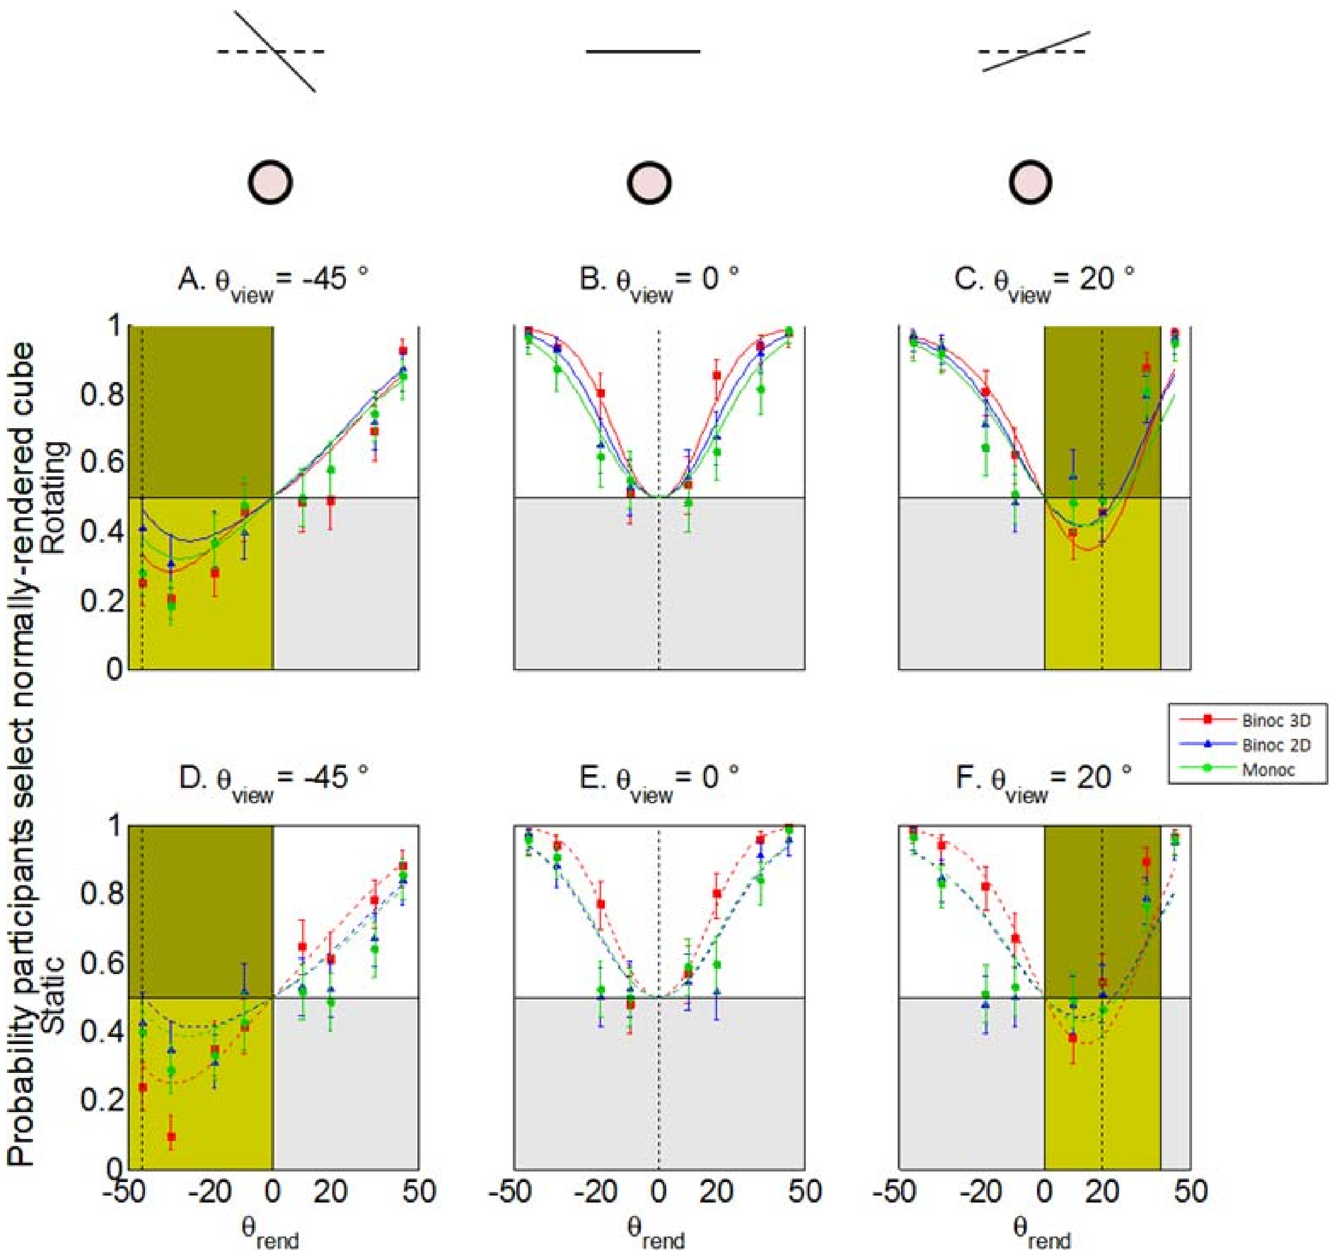


Figure 5. Results from the frame-visible condition. The vertical axis displays the proportion of trials on which subjects reported the normal-rendered cube as appearing "more cube-like." Results are plotted as a function of $\theta_{\text{rend}}$, the viewing angle for which the obliquely rendered cube was drawn. Subjects would therefore necessarily be at chance at $\theta_{\text{rend}} = 0$. Data were pooled across subjects; each data point represents 176 trials from 11 subjects. Error bars show 95% confidence intervals using simple binomial statistics. The top row (solid lines, figures A, B, C) is for trials with rotating cubes, the bottom row (dashed lines, figures D, E, F) is for trials with static cubes. Figures A and D show data from $\theta_{\text{view}} = -45°$, figures B and E from $\theta_{\text{view}} = 0°$, and figures C and F from $\theta_{\text{view}} = 20°$. Shading as in Figure 4.

of the model meant that many different sets of model parameters gave virtually the same value for $\Delta V$ and were thus indistinguishable. To avoid this degeneracy, we set the value of the parameter $A$ to 3 and allowed $B$ to vary. We thus fitted sets of three model parameters ($B$, $r$, $s$) to sets of 24 data points (8 values of $\theta_{\text{rend}} \times 3$ values of $\theta_{\text{view}}$).

## Results

Figures 5 and 6 show the proportion of trials on which the normal-rendered cube was selected as being "more cube-like," pooled over all observers. We plot this as a function of $\theta_{\text{rend}}$, the viewing angle for which

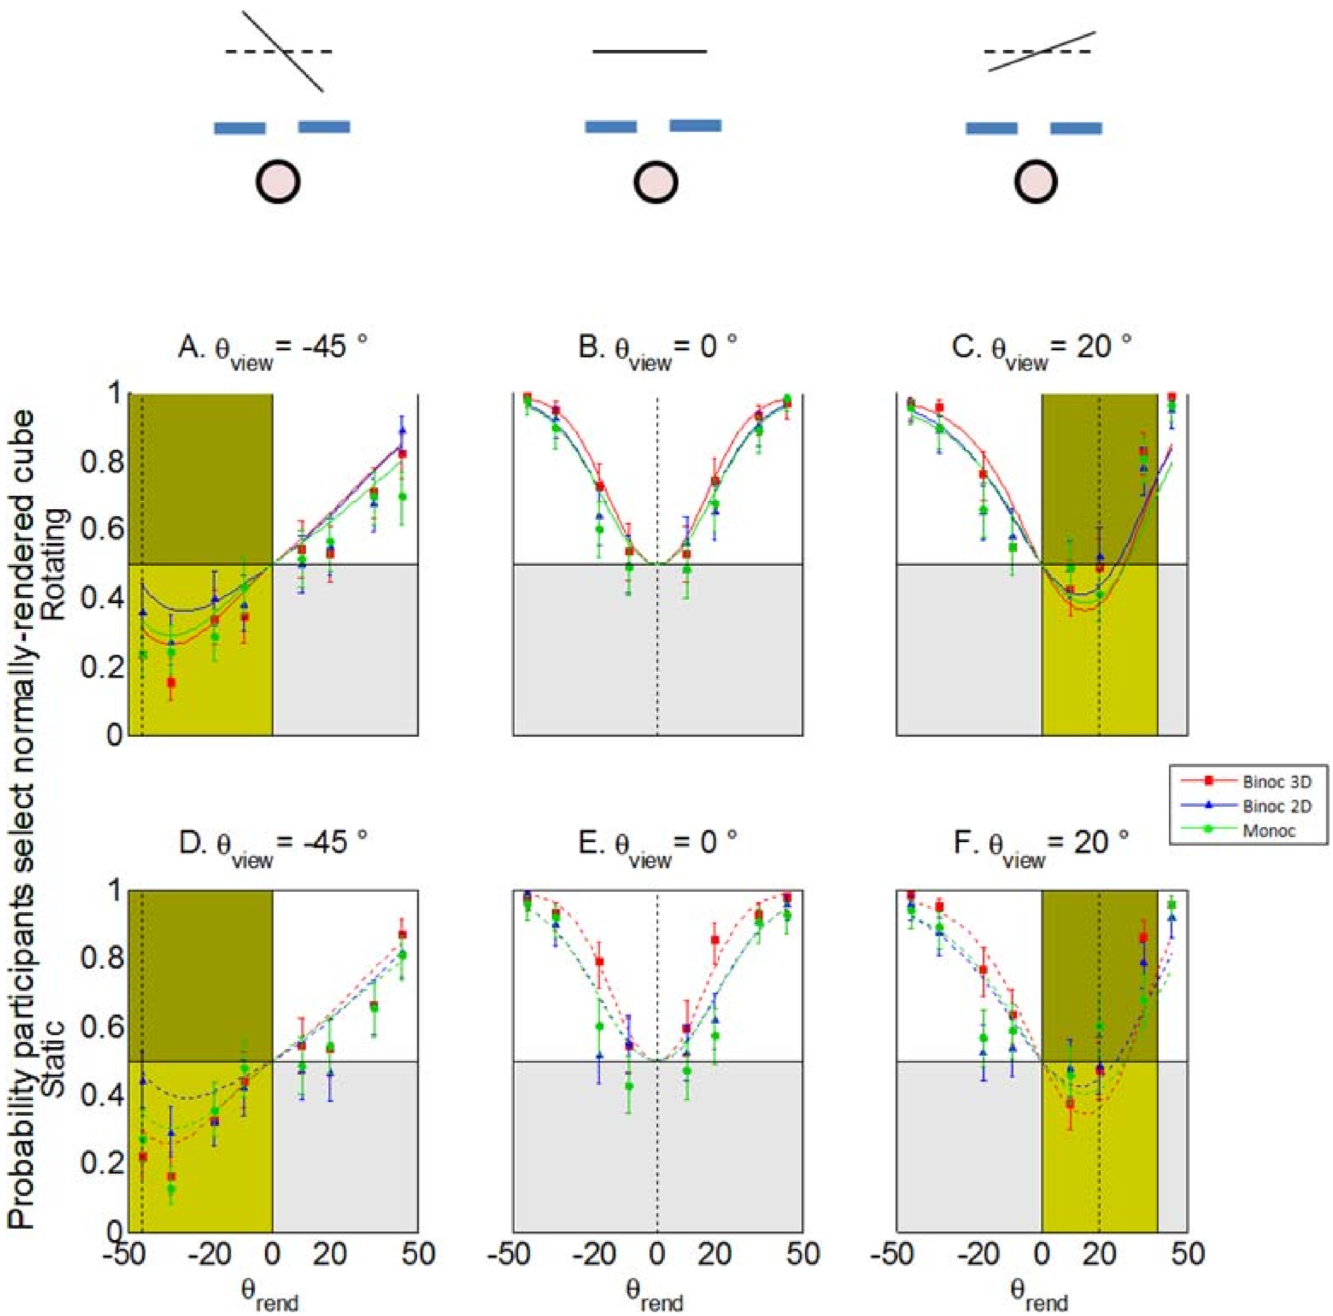


Figure 6. Results from the frame-occluded condition. As for Figure 5, except here a curtain prevented the subjects from seeing the edges of the screen.

the obliquely rendered cube was drawn (Figure 2A). For $\theta_{\mathrm{rend}} = 0$, both cubes would be rendered for perpendicular viewing, so performance would necessarily be at chance. Figures 5 and 6 show results for the frame-visible and frame-occluded conditions, respectively. The three panels in each row show results for the three different viewing angles $\theta_{\mathrm{view}}$. The different colors and symbols show different binocularity conditions: Red squares = binocular viewing in S3D; blue triangles = binocular viewing in 2-D (same image on screen for both left and right eyes); green disks = monocular viewing (pooled left and right monocular results). The upper panels (A through C) show data for rotating stimuli, and the lower (D through F) for static.

Figure 5 shows results for the frame-visible condition, where subjects could see the television screen and thus were aware when they were viewing it obliquely; Figure 6 shows results for the frame-occluded condition, where the edges of the screen were not visible. Above each figure, a schematic is drawn of how the television was orientated and whether a curtain was present or not, as an aid to the reader. The curves in

Figures 5 and 6 show the model fitted to the data as described in the Material and methods. We discuss this in more detail later.

The vertical dashed lines mark the case $\theta_{\text{rend}} = \theta_{\text{view}}$. In this case, for the S3D condition, the obliquely rendered cube should project the same image onto each retina as a real cube (geometrically correct stimulus). The horizontal line at 0.5 marks chance (i.e., both cubes looked equally cube-like to the participant, who thus selected one at random). If objects look veridical when rendered for normal viewing, even when viewed obliquely, data points should lie above this line. If objects look veridical when they are geometrically correct on the retina, where data points should lie depends on rendering and viewing angles. The white regions in each panel show where the normal-rendered cube is closer than the obliquely rendered cube to being geometrically correct for the particular viewing angle. Here the normal-rendered cube should look more veridical, so subjects should select it whenever they can detect a difference between the two render angles (probability $\geq$0.5). The fact that data points do lie in the white regions, rather than in the gray regions below them, confirms this but does not enable us to distinguish between a preference for normal rendering and a preference for geometrical correctness.

Conversely, the yellow regions show where the obliquely rendered cube is closer to geometrically correct. The fact that data points lie predominantly in the bright yellow regions below 0.5 rather than in the dark regions above 0.5 indicates that the preference for geometrical correctness usually won out over that for normal rendering. However, the fact that data points never go as far below 0.5 as above it reveals that viewers were also affected by a preference for normal rendering. This agrees with previous work suggesting that there are two factors which make a virtual object viewed on a screen appear correct to an observer: first, if it creates the same image on the retina as a real object would; and second, if the virtual object *would* create the same image on the retina as a real object *if* the observer were viewing the screen perpendicularly. In the next two sections, we discuss in more detail several aspects of our data which confirm this conclusion.

## Sensitivity to rendering angle $\theta_{\text{rend}}$

We first consider the central panel, Figure 5B, where $\theta_{\text{view}} = 0°$, i.e., the screen was frontoparallel in the usual way. If $\theta_{\text{rend}} = 0°$, both cubes would have the same projection, so performance would be at chance. As the obliquely rendered cube is drawn at ever more extreme angles, it appears progressively more distorted, and subjects become more likely to choose the normal-rendered cube. The rendering angle $\theta_{\text{rend}}$ is significant when considering only this subset of the data ($\chi^2 =$ 42,080.1, $p < 0.0005$). In agreement with previous studies (Cutting, 1987), subjects were fairly insensitive to incorrect rendering. At $|\theta_{\text{rend}}| = 10°$, results do not differ significantly from chance for any binocularity conditions (95% confidence intervals in Figure 5 overlap chance). Even when $\theta_{\text{view}}$ was as large as 20°, the results are not significantly different from chance for a static cube viewed without S3D. For a rotating cube, or a static cube viewed in S3D, subjects were significantly more likely to choose the normal-rendered cube but did so only about 75% of the time. Even when the obliquely rendered cube was drawn for a viewing angle as extreme as 45°, subjects still chose it as being "more cube-like" on nearly 10% of trials when viewing a static cube in 2-D. This is surprising, given that a rendering angle of $\theta_{\text{rend}} = 45°$ produces a very different image on the screen from one of 0° (Figure 3A).

## Effects of oblique viewing angle, $\theta_{\text{view}} \neq 0$

Figure 5A and C shows results where subjects were viewing the screen obliquely. Clearly, the results are very different. At almost every value of $\theta_{\text{rend}}$, participants were less likely to select the normal-rendered cube than when the screen was frontoparallel to them. In the yellow-shaded regions, where a preference for normal rendering conflicts with a preference for geometrical correctness, data points lie in the bright region below chance rather than the shaded region, i.e., participants were more likely to select the object which was closer to geometrically correct. This indicates that they were not able to compensate completely for the oblique viewing angle.

However, oblique viewing clearly had a strong effect on perception, even when the retinal image had been designed to take oblique viewing into account. For example, in Figure 6A, the viewing angle was $\theta_{\text{view}} = -45°$. Thus at $\theta_{\text{rend}} = -45°$, the obliquely rendered cube produced the geometrically correct image of a cube on the retina, whereas the normal-rendered cube was distorted. Figure 6B shows that subjects were quite capable of detecting a 45° error in rendering angle when the display is frontoparallel: They rejected the erroneous rendering over 80% of the time. However, when viewing obliquely at $\theta_{\text{view}} = -45°$ (Figure 6A), subjects did not show a comparably strong preference for the geometrically correct cube: They chose it only 25% of the time for the S3D stimulus at $\theta_{\text{rend}} = -45°$, while for the 2-D stimuli, they picked both cubes equally often. This cannot be explained simply by a lack of sensitivity to distortion (Cutting, 1987; Gombrich, 1972), but must reflect a mechanism favoring normal rendering.

A similar conclusion is indicated by the asymmetry about the line $\theta_{\text{rend}} = \theta_{\text{view}}$ in Figure 5C. Geometrically,

| Factor or interaction | $\chi^2$ | *df* | *p* |
|---|---|---|---|
| Occlusion | 2.723 | 1 | 0.099 |
| Binocularity | 53.290 | 2 | <0.0005 |
| Motion | 6.391 | 1 | 0.011 |
| $\theta_{view}$ | 105.692 | 2 | <0.0005 |
| $\theta_{rend}$ | 4,814.267 | 7 | <0.0005 |
| Occlusion $\times$ Binocularity | 6.461 | 2 | 0.040 |
| Occlusion $\times$ Motion | 1.758 | 1 | 0.185 |
| Occlusion $\times$ $\theta_{view}$ | 50.411 | 2 | <0.0005 |
| Occlusion $\times$ $\theta_{rend}$ | 4,018.044 | 7 | <0.0005 |
| Binocularity $\times$ Motion | 114.703 | 2 | <0.0005 |
| Binocularity $\times$ $\theta_{view}$ | 1,602.966 | 4 | <0.0005 |
| Binocularity $\times$ $\theta_{rend}$ | $>10^{15}$ | 11 | <0.0005 |
| Motion $\times$ $\theta_{view}$ | 10.122 | 2 | 0.006 |
| Motion $\times$ $\theta_{rend}$ | 1,182.191 | 7 | <0.0005 |
| $\theta_{view}$ $\times$ $\theta_{rend}$ | $>10^{14}$ | 10 | <0.0005 |
| Occlusion $\times$ Binocularity $\times$ Motion | 0.624 | 2 | 0.732 |
| Occlusion $\times$ Binocularity $\times$ $\theta_{view}$ | 8.495 | 4 | 0.075 |
| Occlusion $\times$ Binocularity $\times$ $\theta_{rend}$ | $>10^{12}$ | 10 | <0.0005 |
| Occlusion $\times$ Motion $\times$ $\theta_{view}$ | 0.564 | 2 | 0.754 |
| Occlusion $\times$ Motion $\times$ $\theta_{rend}$ | 3,270.183 | 7 | <0.0005 |
| Occlusion $\times$ $\theta_{view}$ $\times$ $\theta_{rend}$ | $>10^{10}$ | 9 | <0.0005 |
| Binocularity $\times$ Motion $\times$ $\theta_{view}$ | 10.000 | 4 | 0.040 |
| Binocularity $\times$ Motion $\times$ $\theta_{rend}$ | $>10^{14}$ | 9 | <0.0005 |
| Binocularity $\times$ $\theta_{view}$ $\times$ $\theta_{rend}$ | 166,597.996 | 8 | <0.0005 |
| Motion $\times$ $\theta_{view}$ $\times$ $\theta_{rend}$ | $>10^{12}$ | 9 | <0.0005 |
| Occlusion $\times$ Binocularity $\times$ Motion $\times$ $\theta_{view}$ | 54.382 | 4 | <0.0005 |
| Occlusion $\times$ Binocularity $\times$ Motion $\times$ $\theta_{rend}$ | 1,768.396 | 8 | <0.0005 |
| Occlusion $\times$ Binocularity $\times$ $\theta_{view}$ $\times$ $\theta_{rend}$ | $>10^{12}$ | 10 | <0.0005 |
| Occlusion $\times$ Motion $\times$ $\theta_{view}$ $\times$ $\theta_{rend}$ | $>10^{12}$ | 9 | <0.0005 |
| Binocularity $\times$ Motion $\times$ $\theta_{view}$ $\times$ $\theta_{rend}$ | $>10^{12}$ | 11 | <0.0005 |
| Occlusion $\times$ Binocularity $\times$ Motion $\times$ $\theta_{view}$ $\times$ $\theta_{rend}$ | $>10^{14}$ | 12 | <0.0005 |

Table 1. Main effects on our results of individual factors and interactions between factors. Results are from a generalized estimating equation done in SPSS that returns the chi-square value along with the degrees of freedom and associated *p* value.

the obliquely rendered cube should appear equally distorted for viewing-angle discrepancies of equal magnitude, $|\theta_{view} - \theta_{rend}|$. Thus, it should appear more distorted for $\theta_{rend} = -10°$ (a discrepancy of 30° from the true viewing angle, $\theta_{view} = 20°$) than for $\theta_{rend} = 35°$ (a discrepancy of only 15°). Yet Figure 5C shows that in fact, for 2-D stimuli, subjects could not perceive the distortion at all for $\theta_{rend} = -10°$ (they picked the obliquely rendered cube as often as the normal-rendered cube), whereas it was fairly obvious to them at $\theta_{rend} = 35°$ (they picked the normal-rendered cube on 75% of trials). This asymmetry, along with the lack of a clear preference for the geometrically correct rendering, is another indication of a compensation mechanism which corrects for oblique viewing and makes objects rendered for normal, perpendicular viewing tend to appear correct even if the retinal image is in fact distorted. However, this compensation works only up to a point. If the compensation were perfect, then Figure 5A and C would be identical to Figure 5B (compare to Figure 4).

## Statistical analysis

Figures 5 and 6 present data with different viewing factors, varying in frame visibility versus occlusion, binocularity, and object motion. We carried out several analyses to assess the effects of these different factors. First, we analyzed the raw data (proportion of normal-rendered selections), which are independent of the assumptions made in our fitted model. We evaluated statistical significance using a generalized estimating equation in SPSS, with intersubject and global comparisons of the raw data that had edge occlusion, object motion, binocularity, angle of projection ($\theta_{rend}$), and viewing angle ($\theta_{view}$) as variables. The five-way interaction yielded significant results ($p < 0.0005$;

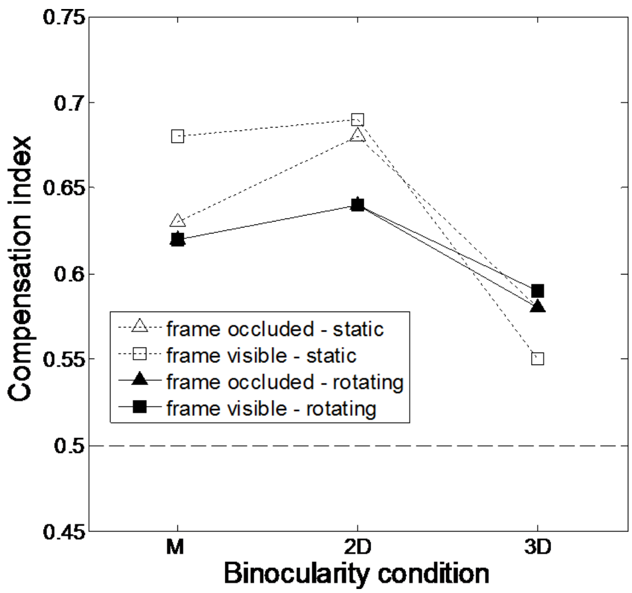


Figure 7. Compensation index $C$, defined as the ratio $A/(A+B)$. $A$ and $B$ are model parameters modeling the strength of the preference for normal rendering and for geometrical correctness, respectively. Higher values of $C$ indicate more compensation for oblique viewing. The dashed horizontal line marks where both weights are equal. This figure shows the values derived from fits to data pooled across all subjects. To carry out the analysis of significance, we derived compensation indices for individual subjects from fits pooled to data from that subject only.

Table 1), but this could be simply due to one specific set of factors yielding a significant result rather than the significance of the factors themselves. Thus we evaluate the main factors and the different possible interactions between the factors in Table 1 at the end of the article.

We discuss the nature and size of these differences in the following sections. The statistical significance of all main effects and interactions are reported in Table 1. We report chi-square values with the degrees of freedom specified.

As can be seen from Table 1, all factors except for edge occlusion had a significant main effect on our results. This implies that, as one would expect, the perceived distortion of the cubes is affected by the angle at which they are viewed and the angle for which they are rendered, as well as by whether they are viewed in S3D, or binocularly or monocularly in 2-D. However, perhaps surprisingly, whether or not the edges of the TV screen are occluded with the curtain does not appear to be important. Most interactions, including all four-way and five-way interactions were also significant.

In the statistical analysis, we considered all the data collected. This makes it difficult to assess the effect of different factors on the two different components identified in our model. As argued earlier, our data imply that two factors affect whether an object appears distorted: whether it is geometrically correct on the retina ($\theta_{\mathrm{rend}} = \theta_{\mathrm{view}}$) and whether it would be correct if viewed perpendicularly ($\theta_{\mathrm{rend}} = 0°$). Much of our data confound these two effects, because often, both factors imply that the user should select the normal-rendered cube. This situation corresponds to the white regions in Figures 6 and 7. To assess how the different experimental conditions (occlusion, binocularity, rotation) affected the competition between the two model components, we also repeated this statistical analysis using only data where the two components pulled in opposite directions, i.e., the yellow regions in Figures 5 and 6. Here there is no overlap in the values of $\theta_{\mathrm{view}}$ and $\theta_{\mathrm{rend}}$, so the statistical significance of $\theta_{\mathrm{view}}$ cannot

| Factor or interaction | $\chi^2$ | *df* | *p* |
|---|---|---|---|
| Occlusion | 32.800 | 1 | <0.0005 |
| Binocularity | 354.095 | 2 | <0.0005 |
| Motion | 5.174 | 1 | 0.023 |
| $\theta_{\mathrm{rend}}$ | 42,080.094 | 6 | <0.0005 |
| Occlusion × Binocularity | 4.183 | 2 | 0.123 |
| Occlusion × Motion | 0.042 | 1 | 0.837 |
| Occlusion × $\theta_{\mathrm{rend}}$ | 35.415 | 6 | <0.0005 |
| Binocularity × Motion | 1.921 | 2 | 0.383 |
| Binocularity × $\theta_{\mathrm{rend}}$ | 816.705 | 8 | <0.0005 |
| Motion × $\theta_{\mathrm{rend}}$ | 155.478 | 6 | <0.0005 |
| Occlusion × Binocularity × Motion | 1.718 | 2 | 0.424 |
| Occlusion × Binocularity × $\theta_{\mathrm{rend}}$ | $>10^{14}$ | 9 | <0.0005 |
| Occlusion × Motion × $\theta_{\mathrm{rend}}$ | 12.864 | 6 | 0.045 |
| Binocularity × Motion × $\theta_{\mathrm{rend}}$ | 228.869 | 8 | <0.0005 |
| Occlusion × Binocularity × Motion × $\theta_{\mathrm{rend}}$ | $>10^{14}$ | 10 | <0.0005 |

Table 2. Effects on our results of individual factors and interactions when the geometrically correct cube was different from the perpendicularly projected cube. Results are from a generalized estimating equation done in SPSS that returns the chi-square value along with the degrees of freedom and associated $p$ value.

be determined. We therefore only consider the main factor influences and the interactions between frame occlusion or visibility, binocularity, rotation, and $\theta_{\mathrm{rend}}$. Table 2 shows the main effects and interaction terms for these four factors.

Within this more limited data set, frame occlusion now had a highly significant main effect on the results, as well as the other factors which did so previously. Considering the two-, three-, and four-way interactions in Table 2, we see that all the interactions including $\theta_{\mathrm{rend}}$ return significant results, whereas any interactions not including $\theta_{\mathrm{rend}}$ are not significant. This makes sense, because clearly the rendering angle $\theta_{\mathrm{rend}}$ is key to whether the object appears distorted. All analysis up to this point is independent of our model. Our statistical analysis implies that frame occlusion, binocularity, and object motion all affect the balance between the competing preferences for a geometrically correct and a normal rendering angle.

## Model fitting

We made this intuitive description quantitative in our two-factor model of perceived veridicality (Equation 1). As Figures 5 and 6 show, it gives a fairly good account of our results. Table 3 gives fitted model parameters and percentage variance explained for the different conditions. The parameters were fitted simultaneously to all data in a given condition for object motion, binocularity, or frame occlusion—i.e., across all three panels in each row, the same parameters are used for all curves of a given color. In every case, the model explains >80% of the variance. Interestingly, the fits are generally better for the binocular S3D and B2D conditions, where they explain >93% and >85% of the variance, respectively, than for the monocular conditions, even though the fit parameters are fitted independently for each binocularity condition. In the monocular conditions, subjects tended to choose the obliquely rendered cube slightly more often than our model can capture, especially when that cube was close to being geometrically correct. However, the generally successful performance of the model confirms the qualitative argument developed earlier, that objects tend to look less distorted if they are rendered *either* for the geometrically correct viewing angle *or* for normal, perpendicular viewing. An advantage of the model is that it also allows us to make quantitative comparisons between the two mechanisms, as follows.

## Quantifying the preference for normal rendering versus geometrical correctness

Our model suggests that the mechanism favoring geometrical correctness is much more sensitive to incorrect rendering angle than that favoring normal rendering. The standard deviations fitted for the Gaussians are 23° and 47°, respectively (means across conditions for data pooled across subjects; Table 1). However, the model suggests that the preference for normal rendering is generally stronger than that for geometrical correctness. The parameter $A$, representing the weight given to normal rendering, is generally larger than $B$, the weight given to geometrically correct images. To quantify this, we define the compensation index (Table 3) as the ratio $C = A/(A + B)$. A value of $C = 0$ would indicate no compensation, such that perception reflected only the geometrical correctness of the image on the retina, without regard for whether the on-screen image would appear correct when viewed normally. A value of $C = 1$ would indicate perfect compensation, such that viewing angle had no effect on perceived veridicality, and no preference for geometrical correctness. Another interpretation of the compensation index becomes apparent when we consider how the perceived veridicality of an object rendered for frontoparallel viewing declines monotonically with viewing angle, relative to its veridicality at frontoparallel viewing. From Equation 1, we have

$$\text{relative veridicality} = \frac{A + B\exp\left(-\frac{{\theta_{\mathrm{view}}}^2}{2r^2}\right)}{A + B} = C + (1 - C)\exp\left(-\frac{{\theta_{\mathrm{view}}}^2}{2r^2}\right). \quad (5)$$

At large viewing angles, this reduces to $C$. Thus, in our model, the compensation index $C$ describes how good a normally rendered picture looks when viewed at the most extreme viewing angles.

Figure 7 plots the compensation index $C$ for the different viewing conditions in our experiment. All 12 data points in Figure 7 lie well above 0.5, indicating that the preference for normal rendering dominates. This may seem surprising, given that in the yellow regions of Figures 5 and 6 where the two preferences conflict, data and model fits both lie below 0.5, i.e., the geometrically correct cube is chosen preferentially. To see why this occurs, it is helpful to consider how the model compares cubes rendered for $\theta_{\mathrm{rend}} = 0°$ (normal) and $\theta_{\mathrm{rend}} = 30°$, when the viewing angle is 45°. To the normal-rendering mechanism (the $A$ term in Equation 1), the normal-rendered cube is perfect and the other cube is less veridical because it is 30° away from the peak of the Gaussian. However, because the Gaussian is broad, the difference is not extreme, so the normal-rendering mechanism has only a weak preference for the normal-rendered cube. Conversely, to the geometric-correctness mechanism (the $B$ term in Equation 1), the obliquely rendered cube looks acceptable—the 15°

| | | Fitted model parameters | | | | |
|---|---|---|---|---|---|---|
| | | Weights | Sensitivity | | | |
| | | Geometrical-rendering weight, $B$ | For geometrically correct rendering, $r$ (°) | For normal rendering, $s$ (°) | Compensation index, $C = A/(A + B)$ | % variance explained |
| Frame-visible (Figure 5) rotating | Monocular | 1.84 | 26.58 | 50.52 | 0.62 | 87.52% |
| | Binocular 2-D | 1.72 | 23.92 | 42.04 | 0.64 | 91.20% |
| | Binocular S3D | 2.12 | 20.59 | 41.81 | 0.59 | 90.98% |
| Static | Monocular | 1.40 | 23.76 | 49.15 | 0.68 | 83.84% |
| | Binocular 2-D | 1.33 | 24.79 | 47.40 | 0.69 | 81.10% |
| | Binocular S3D | 2.47 | 24.00 | 42.16 | 0.55 | 91.31% |
| Frame-occluded (Figure 6) rotating | Monocular | 1.82 | 23.00 | 51.93 | 0.62 | 88.60% |
| | Binocular 2-D | 1.68 | 23.41 | 44.83 | 0.64 | 90.91% |
| | Binocular S3D | 2.18 | 22.65 | 45.01 | 0.58 | 91.23% |
| Static | Monocular | 1.74 | 24.53 | 55.46 | 0.63 | 83.75% |
| | Binocular 2-D | 1.40 | 23.01 | 47.87 | 0.68 | 85.13% |
| | Binocular S3D | 2.20 | 21.12 | 44.01 | 0.58 | 91.91% |

Table 3. Fitted model parameters for weights and sensitivity, including the implied effectiveness of compensation, as well as the percentage variance explained, for all conditions (Equation 2). Table rows are color-coded as in Figures 5 and 6. Note that the normal-rendering weight, parameter $A$, was constrained to be equal to 3, and so is not included in the table (see Material and methods).

error in render angle is less than one standard deviation—but the normal-rendered cube looks very poor, with a 45° error of two standard deviations. This mechanism therefore has a strong preference for the obliquely rendered cube. When the preferences of both mechanisms are summed, the strong preference for the obliquely rendered cube wins out over the weak preference for the normal-rendered cube.

We can use the compensation index to assess how the different viewing factors influence the relative weights of the competing preferences. We have experimented with various metrics, and they all show a similar effect of viewing condition. To evaluate the significance of the different factors, we generated compensation indices for each subject individually and used a linear general estimating equation implemented in SPSS. We found a significant three-way interaction between binocularity, object motion, and frame occlusion ($p < 0.0005$). Most of this appears to come from the significant effect binocularity had on the compensation index. In particular, there was a very significant difference between S3D and both B2D and monocular conditions (both $p$s $< 0.0005$), but there was still a significant difference between B2D and monocular conditions ($p = 0.024$). Frame occlusion ($p = 0.022$) and rotation ($p < 0.0005$) had a significant effect in the monocular condition but not in the other binocularity conditions.

In summary, then, our statistical analyses both of the raw data and of the fitted model parameters imply that binocularity, frame occlusion, and object motion all affect the balance between the preferences for geometrical correctness and normal rendering. In the next sections, we consider each factor in turn.

## Effect of S3D

Figure 7 implies that S3D weakens the compensation mechanism and gives more weight to whether rendered objects create the correct image on the retina. This agrees with the results of Banks et al. (2009). The compensation index drops from $C = 0.66$ for binocular 2-D viewing to 0.575 for stereoscopic S3D viewing (averaged over other viewing conditions): a small but statistically significant difference. This effect can be seen in the raw data, when we compare the S3D results in Figure 5 to the 2-D (red squares vs. blue triangles). This is particularly clear in Figure 5D, where the viewing angle is extreme ($\theta_{\text{view}} = -45°$). When the obliquely rendered cube is close to the correct retinal image ($\theta_{\text{rend}}$ close to $\theta_{\text{view}}$), subjects perceive it as more cube-like than the normal-rendered cube when it is viewed in S3D, and select it >75% of the time. However, when viewed in 2-D, it appears nearly as distorted as the normal-rendered cube and is selected only slightly more than half the time. The effect of S3D is also apparent in Figure 5E, where the screen is viewed perpendicularly. Viewers are more sensitive to errors in rendering angle with S3D than with 2-D or monocular content. In 2-D, a rendering-angle error as large as 20° cannot be distinguished from the correct rendering angle of 0°. In S3D, performance at ±20° is around 75%, suggesting that the error is detected on about half of trials.

The difference between the S3D and other conditions is less pronounced with the solid cubes than with the wire-frame cubes used in our previous study (Hands & Read, 2013). This suggests that a major effect of S3D in that study may simply have been its ability to disambiguate the Necker illusion. Once this illusion is removed through the use of solid cubes, S3D makes less difference to the perceptual compensation mechanisms which lead viewers to select mainly the normal-rendered cube. However, even with solid objects, viewing in S3D does tend to enhance the preference for geometrical correctness.

## Effect of frame occlusion

As our statistical tables show, the effect of frame occlusion was one of the weakest in our statistical analysis. Comparing Figures 5 (frame-visible) and 6 (frame-occluded), little difference is apparent. In Figure 7, the compensation index is barely affected by frame occlusion, moving from $C = 0.63$ when the frame was visible to 0.62 when it was occluded (averaged over other viewing conditions), which is not statistically significant. However, occluding the frame does produce a substantial—and significant—drop in compensation for the monocular static condition (Figure 7). This is in qualitative agreement with results from Vishwanath et al. (2005). These authors found some compensation with monocular viewing when the picture frame was visible, but none for monocular viewing through an aperture. We also saw a significant effect of occlusion when restricting our analysis to data where the normal-rendering and geometrical-correctness preferences make opposite predictions ($p < 0.0005$; Table 2).

The small effect overall of frame occlusion is surprising, since the occlusion did appear to be very effective in removing conscious awareness of screen orientation; even the authors could not reliably say which side of the screen was closer when viewing it through the occluder. Yet this does not seem to have produced a substantial tendency to select the geometrically correct cube over the normal-rendered one. For binocularly viewed stimuli, frame occlusion may have had little effect because disparity and vergence cues to screen orientation remained available to viewers and may have been used unconsciously to compensate for screen slant (Rogers & Bradshaw, 1993; Vishwanath et al., 2005). For monocular stimuli, frame occlusion has more of an effect (Figure 7), and indeed we see a significant difference in the results for $\theta_{\mathrm{view}} = \theta_{\mathrm{rend}} = -45°$. Pooling static and rotating stimuli in Figures 5 and 6, viewers were closer to chance when they could see the screen edges (chose the obliquely rendered cube on 120 out of 352 trials) and preferentially chose the obliquely rendered cube when the edges were occluded (90 out of 352 trials). A similar effect persists at $\theta_{\mathrm{rend}} = -35°$. Elsewhere, the lack of an effect seems to be because our participants were relatively insensitive to the distortions caused by rendering angle and thus did not notice when these distortions were corrected.

## Effect of object motion

Our statistical analysis of the raw data indicates that object motion is a significant factor in both the full data set and the important subset where the preference for geometrical correctness is pitted against the preference for normal rendering ($p = 0.011$ and $p = 0.023$, respectively). However, object motion did not have the effect we expected. We had speculated that structure-from-motion cues might contribute to the compensation mechanism, increasing the preference for normal-rendered objects. In fact, object motion decreased the compensation index for both monocular and binocular 2-D cubes (Figure 7); this was significant in the monocular condition. In stereoscopic 3-D, object motion did tend to increase the compensation index, but the increase was not significant. As Table 1 shows, considering the full set of raw data, there is a significant interaction between object motion and binocularity ($p < 0.0005$). Pairwise comparison shows that object motion has a significant effect even when considering the individual binocularity conditions ($p = 0.011$ for all three conditions of S3D, B2D, and monocular). However, this interaction was not significant when we restricted our analysis to the subset of data in Table 2. We conclude that overall, object motion has little consistent effect on perception.

# Discussion

Still photographs and movies are generally designed to be presented on a surface which is frontoparallel to the viewer. Despite this, they continue to look veridical when viewed from an oblique angle. As discussed in the Introduction, this is partly because we are fairly insensitive to the image distortions produced by oblique viewing, but also because the visual system actively compensates for oblique viewing. This compensation mechanism ensures that an image viewed on a screen is perceived as if the screen were frontoparallel to the observer, even if it is in fact viewed obliquely. Stereoscopic 3-D brings its own complications—e.g., filming with converged camera axes—but has always implicitly relied on the same compensation mechanism previously shown to exist for 2-D displays. However, this assumption has not yet been adequately tested for S3D content.

There are good reasons to imagine that this compensation mechanism might be weaker for stereoscopic 3-D content, mainly because disparity is now not a reliable cue to the location and orientation of the screen plane. Informally, one can experience this by moving from left to right in front of an S3D image. The image appears to move in synchrony with you, as when a portrait's eyes appear to follow the viewer around the room (Koenderink, van Doorn, Kappers, & Todd, 2004; Perkins, 1973) but now extended to the whole depicted object. S3D content is often already affected by a number of distortions, such as the puppet-theater effect or cardboard cutout effect (Banks et al., 2012; Yamanoue, Okui, & Okano, 2006). If oblique viewing produces further distortions in perceived depth or shape, this would be a further problem for creators of S3D content. It would be particularly difficult to address in applications such as 3-D cinema, where content must be viewed by large numbers of people simultaneously.

We examined this issue by comparing images rendered for a range of oblique viewing angles with those rendered for a frontoparallel screen, in both 2-D and S3D. We confirm that the human visual system compensates to some extent for oblique viewing angles (Bereby-Meyer et al., 1999; Perkins, 1973; Vishwanath et al., 2005). Due to this compensation, images tend to appear veridical if they are rendered for normal (orthogonal) viewing, even if actually viewed from an oblique angle. However, we additionally find that a competing factor also affects appearance: Images also tend to appear veridical if they are rendered for the geometrically correct viewing angle. This effect predominates for viewing angles more oblique than about 20°. We have produced a quantitative model which well describes viewers' perceptual judgments on this task across a wide range of viewing and rendering angles.

## Sensitivity to viewing angle

Our results confirm that viewers are relatively insensitive to distortions caused by inappropriate viewing angles. In 2-D, most viewers cannot tell the difference between a stimulus rendered for perpendicular viewing and a stimulus rendered with up to 20° error in viewing angle (Cutting, 1987; Perkins, 1973; see Figure 3B for examples). Our modeling also suggests that viewers are much less sensitive to oblique viewing angle in content that was rendered to be viewed normally than they are to deviations from the geometrically correct viewing angle in content that was rendered for oblique viewing.

## Range over which compensation operates

It seems reasonable to expect that viewers should compensate better for small oblique viewing angles than for large ones. Our data appear to support this. For example, when $\theta_{\text{view}} = 20°$, viewers showed only a weak preference for the geometrically correct cube ($\theta_{\text{rend}} = \theta_{\text{view}}$), suggesting that compensation made the normally rendered cube appear nearly as veridical, whereas when $\theta_{\text{view}} = -45°$, they showed a stronger preference for the geometrically correct cube (Figures 5 and 6, panels A and D vs. C and F). According to our model, pictures appear more veridical for small oblique viewing angles than for large ones (Equation 5). Our model assumes that compensation works equally well for all viewing angles (blue curves in Figure 4). The decline in veridicality comes from the preference for geometrical correctness against which the compensation mechanism is pitted. In our model, taking $C = 0.62$ and $r = 23°$ as representative values, veridicality never drops below 62% of optimal even at the most extreme angles, and remains above 80% even out to viewing angles of 28°.

## Regression to expected shape

Some previous authors have suggested that humans have a tendency to regress distorted images of familiar objects to their expected form (Gombrich, 1972; Thouless, 1931). Presumably, regression is imagined as operating on retinal images to make them appear more geometrically correct. If the regression operated perfectly no matter what the distortion, both objects in our experiment would appear equally cube-like, and performance would be at 50% throughout. The geometric term in our model (the $B$ term in Equation 1) is effectively an implementation of regression which allows for the possibility that regression is more effective for small departures from geometrical correctness. The parameter $r$ describes the range over which regression operates, with perfect regression corresponding to the case $r \to \infty$ and $A = 0$.

## Differences between stereoscopic 3-D and 2-D

In line with expectations, we found that compensation for oblique viewing works better in 2-D images than in stereoscopic 3D. A plausible reason is that, in binocular 2-D viewing, the true orientation of the screen can be deduced from binocular cues such as disparity or vergence, even when the edges of the screen are occluded from view. This makes it possible to apply the appropriate compensation (Vishwanath et al., 2005).

Banks et al. (2009) also compared shape distortions in oblique viewing for 2-D and S3D stimuli and found that viewing in S3D abolished compensation almost completely. In contrast, we find that compensation still dominates even in S3D, though it is less effective than in 2-D. We highlight three differences in protocol which may contribute to this difference. First, perceptual invariance depends on the stimulus, and particularly the depth variation in the stimulus (Banks et al., 2009). Our stimuli were solid cubes, with a side length from 6 to 14 cm, viewed from a distance of 120 cm. The stimuli of Banks et al. were hinged wire-frame squares with a side length of 30 cm, viewed from 45 cm. Our stimuli thus contained relatively less depth variation. Second, the longer viewing distance used in our study may have enhanced the preference for normal rendering. Artists since the Renaissance have discussed the recommended distance at which to capture the perspective projection of an object in order for it to look pleasing and natural. Hagen and colleagues have argued for a distance at least 10 times the mean object size along its various dimensions, very close to that used in our experiments (Hagen & Elliott, 1976; Hagen et al., 1978); Leonardo da Vinci recommended a smaller distance of 3 times the height of the object (Da Vinci, 2012). The longer viewing distance reduces the amount of perspective convergence, making the projection closer to orthographic. Viewers report such projections as appearing more veridical, even when they are geometrically incorrect for the given viewing distance (Hagen & Elliott, 1976; Hagen et al., 1978). This reflects the fact that viewers do not compensate for wrong viewing distance as they do for oblique viewing (Cooper, Piazza, & Banks, 2012). Thus, when viewed and rendered normally, our cubes should have looked veridical, whereas the hinge stimuli of Banks et al. may still have looked wrong because of the short viewing distance relative to the size of the object. This may have weakened the effectiveness of the compensation for oblique viewing.

Finally, our stimuli were renderings of cubes, where there is a clear canonical form which may have influenced perception, whereas those of Banks et al. were wire-frame hinges, with no clear expectation regarding hinge angle. One might expect the regression mechanism of Thouless (1931) and Gombrich (1972) to operate more strongly on cubes than on hinges. In the terms of our model, this would be expected to boost the parameter $r$, i.e., make subjects more tolerant of departures from geometrical correctness. It might also boost the weight of $B$ relative to $A$, thus reducing the compensation index $C$. If so, this could potentially be one reason we found less compensation with cubes than Banks et al. did with hinges.

Our longer viewing distance and use of familiar objects makes our study more relevant to typical applications of S3D displays in entertainment. The S3D entertainment industry can therefore be reassured by the lack of difference we found between S3D and 2-D content, even at a viewing angle as large as 20°. The differences only became apparent at the largest viewing angle used, 45°. For most S3D display systems, such an extreme viewing angle already causes other problems, such as increased cross talk or contrast changes.

## Effect of frame visibility and object motion

A new contribution of our study was that we investigated the effect of object motion. This is particularly relevant for entertainment applications of S3D, where content is generally dynamic. We had speculated that structure-from-motion cues, together with the rigidity heuristic, might enable the visual system to compensate more effectively for oblique viewing. In fact, object motion had little effect in S3D and tended to weaken compensation in the 2D and monocular conditions. This suggests that the difference between S3D and 2-D television and movies may be even less than that for S3D and 2-D static images.

## Limitations

Our experiment suffered from high levels of cross talk. Cross talk (or ghosting) refers to any leaking of the left eye's image into the right eye and the right eye's images into the left. This can disrupt the perceived depth of the image and lead to double vision, as both eyes see some part of both the stereoscopic images displayed on the screen. High levels of cross talk can lead to incorrect perception of the image seen and affect image quality, so it is essential to minimize cross talk to achieve high-impact, impressive 3-D images (Woods, 2011). Since the 3-D television was manufactured for perpendicular viewing, the amount of cross talk between the images increased substantially with oblique viewing: up to 7%. The fact that the experiments were conducted in darkness also tended to make any cross talk more visible to observers. Cross talk could mean that our "monocular" images are in fact 2-D binocular images in which one image has much lower contrast than the other. Thus, our experiments may underestimate the difference between the 2-D and monocular stimuli, especially for oblique viewing. However, we did repeat some of the monocular conditions with one eye covered, instead of using the 3-D glasses to interleave monocular and binocular stimuli, and we obtained broadly similar results. When it comes to S3D TV, the cross talk increases the ecological relevance of our study, since these levels of cross talk are those which would be experienced by S3D TV viewers in a normal home environment.

The condition of monocular viewing when the frame was occluded from view (green data in Figure 6A, C, D, and F) was intended to remove all information about screen orientation, by removing disparity and perspective cues. However, we were evidently not successful in this, since there was evidence of active compensation even in this monocular, frame-occluded condition. For example, at a viewing angle of $\theta_{\text{view}} = 20°$, our data show an asymmetry in the effect of render angle, indicating that objects looked more cube-like when rendered for a viewing angle closer to frontoparallel than the actual viewing angle than they did when the render angle was equally distant from the actual viewing angle but in the opposite direction. This must mean that subjects had access to some source of information about screen orientation. Possible sources of information include accommodation, motion parallax from small head movements within the headrest, gradients in luminance across the screen, and so on. However, this limitation does not affect our main conclusion, which relates to the difference between binocular 2-D and S3D viewing. Less surprisingly, in this impoverished viewing condition, subjects had greater uncertainty and were less able to perceive any differences between the two cubes. Our model fits indicate lower sensitivity under monocular viewing in almost all cases.

In debriefing after the experiment, several participants stated that they tended to choose the smaller cube when the task was difficult, presumably because any deviations from the canonical cubic form are harder to detect in smaller objects. Since the size of our cubes was chosen at random, this strategy would push performance towards chance, making it harder for us to detect effects of our experimental parameters.

This study has only considered one effect of oblique viewing: distortions in perceived shape. Another approach would be to consider whether viewing stereoscopic content from inappropriate viewing angles is a source of viewer discomfort (Howarth, 2011; Lambooij, Ijsselsteijn, Fortuin, & Heynderickx, 2009). It would be interesting to look at the various different definitions of a zone of comfort, the range of depth allowed in 3-D content before discomfort begins to adversely affect the viewing, and see if changing the viewing angle has any effect on the zone of comfort (Shibata, Kim, Hoffman, & Banks, 2011).

Finally, our study only asked viewers to consider which of two objects most resembled a perfect cube. We did not assess what the viewers were using to make this distinction, nor how they perceived the objects. Accordingly, our model also only predicts perceptual judgments in this comparison task, rather than directly predicting perceived shape.

## Conclusion

When viewing a familiar object, especially one in motion, viewers are very nearly as tolerant to oblique viewing in S3D as in 2-D. This is partly because viewers are fairly insensitive to incorrect viewing angle and partly because of a compensation mechanism which makes content rendered or filmed for a frontoparallel screen continue to appear veridical even when viewed obliquely. Contrary to previous literature suggesting that this compensation is substantially impaired for S3D content, we find little difference. This helps explain why S3D content is popular and effective even though it is usually viewed from the "wrong" position.

*Keywords: stereoscopy, binocular vision, oblique viewing, correction*

## Acknowledgments

This study was partly funded by BSkyB and the Engineering and Physical Sciences Research Council as a Collaborative Awards in Science and Engineering PhD studentship supporting PH. Thanks to Suzanne Pinkney for running the volunteer recruitment scheme. All data presented here were collected by PH for use in this study.

Commercial relationships: BSkyB (PH).
Corresponding author: Paul Hands.
Email: paul.hands@ncl.ac.uk.
Address: Institute of Neuroscience, Newcastle University, Tyne and Wear, UK.

## References

Banks, M. S., Held, R. T., & Girshick, A. R. (2009). Perception of 3-D layout in stereo displays. *Information Display (1975), 25*(1), 12–16.

Banks, M. S., Read, J. C. A., Allison, R. S., & Watt, S. J. (2012). Stereoscopy and the human visual system. *SMPTE Motion Imaging Journal, 121*(4), 24–43.

Bereby-Meyer, Y., Leiser, D., & Meyer, J. (1999). Perception of artificial stereoscopic stimuli from an incorrect viewing point. *Perception & Psychophysics, 61*(8), 1555–1563.

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10*(4), 433–436.

Cooper, E. A., Piazza, E. A., & Banks, M. S. (2012). The perceptual basis of common photographic practice. *Journal of Vision, 12*(5):8, 1–14, http://

www.journalofvision.org/content/12/5/8, doi:10.1167/12.5.8. [PubMed] [Article]

Cutting, J. E. (1987). Rigidity in cinema seen from the front row, side aisle. *Journal of Experimental Psychology: Human Perception & Performance, 13*(3), 323–334.

Da Vinci, L. (2012). *A treatise on painting*. Mineola, NY: Courier Dover.

De La Gournerie, J. (1859). *Traité de perspective linéaire contenant les tracés pour les tableaux, plans & courbes, les bas reliefs & les décorations théatrales, avec une théorie des effets de perspective* [Translation: Discusses Linear Perspective: Containing Traces To Tables, Maps And Curves]. Dalmont et Dunod.

Deregowski, J. B. (1969). A pictorial perception paradox. *Acta Psychol (Amst), 31*(4), 365–374.

Dodgson, N. A. (2004, January). *Variation and extrema of human interpupillary distance*. Paper presented at Stereoscopic Displays and Virtual Reality Systems XI, San Jose, CA.

Gombrich, E. H. (1972). *Art and illusion: A study in the psychology of pictorial representation* (4th ed.). London, UK: Phaldon.

Hagen, M. A., & Elliott, H. B. (1976). An investigation of the relationship between viewing condition and preference for true and modified linear perspective and adults. *Journal of Experimental Psychology: Human Perception & Performance, 2*(4), 479–490.

Hagen, M. A., Elliott, H. B., & Jones, R. K. (1978). A distinctive characteristic of pictorial perception: The zoom effect. *Perception, 7*(6), 625–633.

Hands, P., & Read, J. C. (2013). *Perceptual compensation mechanisms when viewing stereoscopic 3D from an oblique angle*. Paper presented at the 2013 International Conference on 3D Imaging, Liege, Belgium.

Held, R. T., & Banks, M. S. (2008). Misperceptions in stereoscopic displays: A vision science perspective. *ACM Transactions on Graphics, 2008,* 23–32, doi:10.1145/1394281.1394285.

Howarth, P. A. (2011). Potential hazards of viewing 3D stereoscopic television and cinema: A review. *Ophthalmic & Physiological Optics, 31,* 111–122.

Jahoda, G., & McGurk, H. (1974a). Development of pictorial depth perception: Cross-cultural replications. *Child Development, 45*(4), 1042–1047.

Jahoda, G., & McGurk, H. (1974b). Pictorial depth perception: A developmental study. *British Journal of Psychology, 65*(1), 141–149.

Kleiner, M., Brainard, D., & Pelli, D. (2007). *What's new in Psychtoolbox-3?* Paper presented at the 30th European Conference on Visual Perception, Arezzo, Italy.

Koenderink, J. J., van Doorn, A. J., Kappers, A. M. L., & Todd, J. T. (2004). Pointing out of the picture. *Perception, 33,* 513–530, doi:10.1068/p3454.

Kubovy, M. (1988). *The psychology of perspective and Renaissance art*. Cambridge, UK: Cambridge University Press.

Kurtz, H. F. (1937). Orthostereoscopy. *Journal of the Optical Society of America, 27*(10), 323–339.

Lambooij, M., Ijsselsteijn, W., Fortuin, M., & Heynderickx, I. (2009). Visual discomfort and visual fatigue of stereoscopic displays: A review. *Journal of Imaging Science and Technology, 53*(3), 1–14.

Mayhew, J. E., & Longuet-Higgins, H. C. (1982). A computational model of binocular depth perception. *Nature, 297*(5865), 376–378.

Necker, L. A. (1832). Observations on some remarkable optical phaenomena seen in Switzerland; and on an optical phaenomenon which occurs on viewing a figure of a crystal or geometrical solid. *London and Edinburgh Philosophical Magazine and Journal of Science, 1*(5), 329–337.

Ogle, K. N. (1938). Induced size effect I: A new phenomenon in binocular vision associated with the relative size of the images in the two eyes. *Archives of Ophthalmology, 20,* 604–623.

Olson, R. K., & Boswell, S. L. (1976). Pictorial depth sensitivity in two-year-old children. *Child Development, 47*(4), 1175–1178.

Pelli, D. G. (1997). The Video Toolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10*(4), 437–442.

Perkins, D. N. (1973). Compensating for distortion in viewing pictures obliquely. *Perception & Psychophysics, 14*(1), 13–18, doi:10.3758/Bf03198608.

Pirenne, M. H. (1970). *Optics, painting & photography*. London, UK: Cambridge University Press.

Read, J. C., Phillipson, G. P., & Glennerster, A. (2009). Latitude and longitude vertical disparities. *Journal of Vision, 9*(13):11, 1–37, http://www.journalofvision.org/content/9/13/11, doi:10.1167/9.13.11. [PubMed] [Article]

Rogers, B. J., & Bradshaw, M. F. (1993). Vertical disparities, differential perspective and binocular stereopsis. *Nature, 361*(6409), 253–255.

Rosinski, R. R., Mulholland, T., Degelman, D., & Farber, J. (1980). Picture perception: An analysis of visual compensation. *Perception & Psychophysics, 28*(6), 521–526.

Saunders, J. A., & Backus, B. T. (2007). Both parallelism and orthogonality are used to perceive

3D slant of rectangles from 2D images. *Journal of Vision, 7*(6):7, 1–11, http://www.journalofvision.org/content/7/6/7, doi:10.1167/7.6.7. [PubMed] [Article]

Shibata, T., Kim, J., Hoffman, D. M., & Banks, M. S. (2011). The zone of comfort: Predicting visual discomfort with stereo displays. *Journal of Vision, 11*(8):11, 1–29, http://www.journalofvision.org/content/11/8/11, doi:10.1167/11.8.11. [PubMed] [Article]

Thouless, R. H. (1931). Phenomenal regression to the real object. *Nature, 131,* 261–263.

Ullman, S. (1979). The interpretation of structure from motion. *Proceedings of the Royal Society of London B: Biological Sciences, 203*(1153), 405–426.

Vishwanath, D., Girshick, A. R., & Banks, M. S. (2005). Why pictures look right when viewed from the wrong place. *Nature Neuroscience, 8*(10), 1401–1410.

Wheatstone, C. (1838). Contributions to the physiology of vision. Part the first. On some remarkable, and hitherto unobserved, phenomena of binocular vision. *Philosophical Transactions of the Royal Society of London, 128,* 371–394.

Woods, A. J. (1993, September). *Image distortions in stereoscopic video systems*. Paper presented at Stereoscopic Displays and Applications IV, San Jose, CA.

Woods, A. J. (2011). *How are crosstalk and ghosting defined in the stereoscopic literature?*. Paper presented at Stereoscopic Displays and Applications XXII, San Francisco, CA, USA.

Woods, A. J., Docherty, T., & Koch, R. (1993). Image distortions in stereoscopic video systems. *Stereoscopic Displays and Applications IV, 1915,* 36–48, doi:10.1117/12.157041.

Yamanoue, H., Okui, M., & Okano, F. (2006). Geometrical analysis of puppet-theater and cardboard effects in stereoscopic HDTV images. *IEEE Transactions on Circuits and Systems for Video Technology, 16*(6), 744–752, doi:10.1109/TCSVT.2006.875213.

Zorin, D., & Barr, A. H. (1995). Correction of geometric perceptual distortions in pictures. Paper presented at the Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques, Los Angeles, CA, USA.