

This is the author's version. It may currently be cited as:

Durgin, F. H., Ruff, A. J., & Russell, R. (in press). Constant enough: On the kinds of perceptual constancy worth having. In G. Hatfield and S. Allred (eds.), *Visual Experience*. Oxford University Press.

## **Constant enough: On the kinds of perceptual constancy worth having**

Frank H. Durgin, Anna J. Ruff, and Robert C. Russell

Swarthmore College

Perceptual constancy • Visual experience • Slant perception • Monocular vision • Binocular vision • Demand characteristics • Introspection • Perceptual information • Perceptual processing • Space perception

Correspondence to: Frank H. Durgin, Swarthmore College, Department of Psychology, 500 College Ave, Swarthmore, PA 19081, fdurgin1@swarthmore.edu, +01 610-328-8678

### **Abstract**

*The idea is proposed that whereas perceptual experience is underconstant in one sense, it is virtually constant insofar as it is functionally stable and predictable. The possibility of distinguishing perception and cognition is explored in experiments on the perception of surface orientation. These experiments are related to the study of self-motion perception and space perception. An experiment comparing monocular and binocular perception of hills, revealed perceptual differences, between-subjects, that were masked in within-subject comparisons by metacognitive strategies. A second experiment found that participants wearing heavy backpacks gave (cognitively) elevated slope estimates only because of experimental demands not physical ones. Perceptual experience is informative about perceptual processing, but reports of experience are subject to cognitive contamination. True perceptual experience may be virtually constant insofar as the perceptual consequences of actions can be correctly anticipated.*

## Introduction

When humans walk through a fixed environment they experience that environment as stationary and stable. Yet at the same time they may be aware of the optic flow of the objects and walls around them. This awareness of motion superimposed on a stable world might be compared to the awareness of shadows or other gradients of light cast upon a surface that nonetheless appears uniform in color: We know that we are the source of the motion, just as we know that illumination changes are the sources of shadows. But awareness of optic flow might alternatively be compared to the awareness of converging lines in a hallway whose walls nonetheless appear parallel. We do know that the perceived optic flow rate during self-motion is much reduced compared to how the same relative flow rate would be experienced when standing still (Durgin *et al.* 2005). That is, "retinal" properties of flow are lost during self-motion (like "retinal" size and color are unrecoverable). Thus, our perceptual experience seems to include only the shadows of underconstancy, which may often be corrected for in judgment by metacognitive awareness (Granrud, 2009).

A primary thesis of this paper is that this state of affairs (systematic underconstancy) may be no accident if our perceptual experience needs to retain the structured correlations among perceptual and motor variables that can most effectively control and guide our actions. Whereas Noë (2004) suggests that perceptions of constancies, like surface color, emerge, Gibson-like, from the external predictability of the non-constant image transformations, he seems to assume that we have access to the undistorted retinal array. An alternative view is that we have partial constancy in our immediate perceptual experience and that virtual constancy and stability can be achieved insofar as the perceptual predictions we can make about the consequences of our actions are accurate. That is, virtual perceptual constancy is achieved insofar as we can predict how our (non-constant) perceptions will be modified as we move our bodies (e.g. down a hallway) or even as we move our attention.

Full perceptual constancy, on this account (e.g. not being able to see optic flow at all) might be disastrous. For example, when we move, objects that we pass rotate with respect to our reference frame, but Wallach *et al.* (1974) showed that during self-motion there is a huge drop in sensitivity to correlated object rotations (relative to an absolute reference frame). This is probably because precise predictions of object shear for stationary objects in the environment during self-motion would require a full and accurate coding of spatial layout, so it is probably more efficient to accept nearly whatever shear emerges during self-motion as consistent with a stable environment (discrimination is sacrificed for stability).

In contrast, Durgin and Gigone (2007) have shown that sensitivity to optic flow rates for highly predictable surfaces like the ground plane is enhanced during self-motion (for speeds in the range appropriate to walking, for example), consistent with the idea that predictability is fundamental: Retinal flow is distorted (consistent with partial constancy in the form of world stability), but the distortion produces a gain in sensitivity to information relevant to the prediction of perceptual consequences of self-motion (Durgin 2009; see also Abrams *et al.* 2007). Partial constancy may thus provide a useful compromise.

A secondary thesis of this paper however, is that perceptual experience, however tainted and distorted, has a kind of stability and reliability that can seem to be falsified by incautious methods of measurement. Whereas the boundary between perception and judgment is never easy to define, there are contexts in which it is possible to distinguish between judgmental biases and perceptual ones. Finding ways to make these distinctions with greater certainty may not be easy but is important. Although the study of

perceptual processing cannot be conducted reliably by introspection alone, neither should it be conducted as if introspection were impossible. The studies of perceived optic flow rate discussed above depended on several different forms of measurement including, importantly, magnitude estimation, which is a form of momentary introspection. But any particular measurement technique can introduce biases and confusions. For example, successive comparison of two stimuli may appear to be a fairly safe perceptual task, but judgments relative to an internal standard have been shown to be much more precise than perceptual comparisons of successively presented stimuli (Nachmias 2006). Even successive comparison involves memory and therefore the possibility of judgmental bias.

Here we consider some problems in the empirical conceptualization of perceptual experience using the perception of empty spatial extents and two studies of surface orientation as sample cases. We start with the idea that perceptual experience is not to be confused with perceptual information and then develop the idea that neither it is to be confused with perceptual judgments. We will return at the end to the possible relation between perceptual underconstancies and the kind of virtual perceptual constancy that successful perceptual prediction affords.

### **One depth doesn't fit all**

One preliminary point to be made is that different sources of perceptual information may lead to different kinds of perceptual experience and that constancy of the sort we are discussing may fail when certain specific forms of information are absent. Some theories of cue combination suggest that when *visual* sources of information about depth can be integrated, they will be (e.g. Hillis *et al.* 2002), whereas cross-modal sensory experiences are not necessarily fused. Hillis *et al.* considered the fusion of stereoscopic information about surface slant with texture information about surface slant and contrasted it with haptic information about surface slant. There are, of course, crossmodal confusions. For example, many individuals report *feeling* the warmth of a bright laser pointer light shone on "their" hand (even though they are only viewing a light shown on a rubber hand that, by means of mirrors, appears in the location of their hand; Durgin *et al.* 2007a). In speech, most individuals report *hearing* "da" when seeing video of clearly separated lips pronouncing the syllable "ga" in synchrony with a sound that normally sounds like "ba" (McGurk and MacDonald 1976). If one assumes that perception is an *attribution* based on available evidence, it seems reasonable to assume that depth information provided by one rich (metric) source of visual information can be substituted for that produced by another. However, there appear to be strange incompatibilities between different sources of information about relative distance from the observer. Indeed, some of these incompatibilities lie within a given system.

For example, Mamassian (2008) has found that perceiving that a surface is slanted using binocular disparity is 10 times harder (less sensitive) than perceiving a difference in depth between two surfaces – even though the stimulus information specifying the presence of slant and the difference in depth was identical in his stimuli. Perceived slant is simply a different feature than distance, which receives different information from disparity maps.

Whereas textbooks commonly suggest that motion parallax is a rich source of metric depth information (e.g. Wolfe *et al.* 2006), and there are even arguments that it is preferable on computational grounds (e.g. Richards 1985), the case of Stereo Sue (Sacks 2006) provides a hint that this is not reflected in perceptual experience. Susan Barry was born strabismic (cross-eyed) so that her eyes did not focus on the same place. Her vision in each eye was perfectly adequate, and her visual experience for most

of her life alternated between her two eye views, with each eye image sometimes being the suppressed one. Because she never used her eyes in concert, she was stereoblind until in middle age she undertook vergence exercises that allowed her to (eventually), stop suppressing and use her two (now verged) eye images in conjunction with one another. Her case is already highly controversial, because of the widespread belief that unused stereo channels atrophy completely. Moreover, her subjective reports concerning the experiential difference between her new binocular perceptions and her former experience of the world has only increased the level of skepticism about her claims. This is because she suggested that prior to gaining binocular vision, she had never experienced "depth", by which she may mean the empty space between things. For most current theorists, this notion seems confused. Depth is depth, they argue, by whatever source. Such theorists accept that different sources of depth information may be more precise than others, but assume that all provide perceptual "depth".

One specific example Sue gives is of seeing for the first time that trees contain a tangible volume of space (as opposed to merely a layered tangle of branches). The tree example seems telling because it is exactly the example used by Wolfe *et al.*'s (2006) excellent textbook when describing the richness of depth provided by motion parallax. Their textbook suggests that the student lie under a tree...

Gaze up in the branches and leaves with one eye covered and your head stationary. You will notice that the leaves and branches form a relatively flat texture. You can see all the details, but you may have trouble deciding if one little branch lies in front of or behind another. If you open the other eye, stereopsis ... will allow the branches and leaves to fill out a three-dimensional volume that was lacking before. Close the eye and the volume collapses. Now, move your head from side to side and motion parallax will restore the sense of depth. (Wolfe *et al.* 2006, 137).

Although Wolfe *et al.* (2006) intend to portray motion parallax as equivalent to stereopsis, the choice of words here creates an appropriate contrast between them insofar as "the *sense* of depth" one gets from motion parallax sounds entirely inferior to the "three-dimensional volume" provided by stereopsis. LeClair and Durgin (2008) compared metric depth interval estimation from motion parallax with that from binocular stereopsis. We suspended pairs of objects ("clouds" of polyester batting) at different locations within a well-lit, but featureless chamber four meters deep and asked observers to estimate the true distances between the paired objects in one of three conditions:

Monocularly, with lateral head movements of twice the typical interpupillary distance (motion parallax).

Binocularly, without head movements, (binocular stereopsis), or

Monocularly, without head movements, (control)

The difference in subjective impressions for motion parallax and binocular conditions, consistent with the retrospective claims of Susan Barry, were reflected by depth interval estimates which were much less variable and much more accurate under static binocular viewing than with motion parallax. Indeed, although motion parallax reduced response variability relative to the control condition, the average estimates in the two monocular conditions were quite similar and were much less than the depth estimates given under binocular viewing. Our own subjective impression was that as we moved our head with only one eye open we saw the two objects slide back and forth with respect to each other, but we (like our naïve participants, apparently) experienced

no sense of depth between them, only the sure knowledge that larger relative motions signaled larger separations in depth.

Our observation is not without precedent. Ono *et al.* (1986) have noted that only in very near space of 30-80 cm do simulated motion parallax displays appear rigid (see also Nawrot 2003). But the more widely held view is the one presented in the textbook, that motion parallax provides all the information you need to get metric depth. We do not deny that some motion parallax displays are incredibly compelling (especially those involving very near space), but we suspect that for normal large scale scenes, motion parallax can often seem impressively effective at capturing spatial layout primarily because the motion parallax gradient can be anchored to the ground plane – which has its own distance metric built into it (Beusmans 1998; Gibson 1950).

So does this support Stereo Sue's assertion that depth from binocular disparity is qualitatively different than depth from other sources? Perhaps, but not convincingly. It remains reasonable, as in the cases examined by Mamassian (2008), to consider that the computations carried out on various kinds of theoretically useful information may lead to surprisingly divergent consequences. The fact that performance at our cloud task was so poor with motion parallax relative to performance with binocular stereopsis encourages the view (consistent with Mamassian's observation) that binocular information is particularly good at representing the volumetric separations between surfaces whereas motion parallax is not particularly good at this – at a viewing distance exceeding a meter. But the striking failure of binocular slant perception in Mamassian's special case is worth keeping in mind. It may be that even though spatial layout was evident to Susan Barry before she became Stereo Sue, there really was no perception of empty space. It was Gibson who suggested that there was no such thing as space perception – only the perception of surfaces. But the cloud study points to empty-space perception as an experiential reality that binocular vision may uniquely support but that is frequently supplemented by metacognition.

It remains possible that the vergence system itself plays an important role in this, but the main point is that the specific kind of volumetric experience that Wolfe *et al.* (2006) describe for the binocular view of trees (which goes beyond what vergence could possibly achieve) really may be something that can only be activated by binocular inputs or by other inputs (e.g., motion parallax in very near space) that have become appropriately linked to binocular inputs through experience. Sometimes different sources of information about surface layout may be overlaid rather than fused. We just don't know enough yet to be sure. What we do know is that perceptual experience is not the same as perceptual information, and yet perceptual experience may have a lot to tell us about how perceptual information is processed. Getting access to perceptual experience for objective measurement is not easy however.

### **Experiment 1: When big effects have small consequences**

One sub-goal in the experimental study of perceptual experience is for the experimenter to convince research participants to simply report their experience honestly. This should help make modeling perceptual experience possible. Some of us develop complicated instruction sets designed to encourage participants to appreciate that their reports are our only access to their subjective experiences and emphasizing that it is those experiences that we want to measure. But there are times (unlike the cloud case above) when the perceptual reports that we collect are in striking variance with our own subjective impressions of how the stimulus appears. These moments can be quite frustrating, indeed, and all the more so because our participants are sometimes a little too clear that despite their original promises to report things as they saw them,

they could not overcome the desire to be right, the wish not to appear foolish. Granrud's (2009) documentation of meta-perceptual awareness in children engaged in size estimation tasks is a striking example of how dual awareness is one of the facts of perceptual reports.

One of the fundamental principles of slope perception is known as the frontal tendency (Gibson 1950) – the tendency of surfaces to appear steeper than they are. Bridgeman and Hoover (2008) have recently demonstrated, for example, that farther portions of hills appear steeper than nearer portions – arguing that this is partly because the visual information available to see the slopes as departing from vertical becomes weaker with distance. We conducted an experiment seeking to demonstrate that a fairly steep hill (of about 20 deg) would appear steeper when viewed monocularly than when viewed binocularly. We had observed the effect ourselves. The effect is well known for small texture-defined surfaces viewed through an aperture (e.g. Gibson and Cornsweet, 1952), and we sought to document it for large locomotor surfaces. When we closed an eye and looked straight ahead at the 20-deg hill it appeared to us about 60-70 deg. With both eyes open it looked about 40 deg to us (slope overestimation is a typical and persistent finding).

The participants for our study met us near a campus field house. They were then blindfolded and led to one of two grass-covered hills. There they stood between two barriers that blocked any side view of the hill and were allowed to look straight at the hill with either one or two eyes. A cluster of small white stones placed on the hill at approximately eye level served as a fixation mark. Our instructions were clear. We made them read a statement explaining that we wanted them to tell us how things looked, not how they believed them to be. We went over the instruction again orally – we emphasized the importance of them reporting their perceptual experience. We even had them use a palm board first – a board that is suspended on a horizontal axis and can be set to any orientation – because using palm boards avoids some of the numeric bias effects that go with verbal estimates. The (unseen) palm board estimate having been recorded, each participant then made a verbal judgment of the apparent slope of the hill. We then blindfolded them again, led them to the other of the two hills and had them make the same pair of judgments (palm-board match and verbal estimate) in the other viewing condition (monocular if they were binocular at first, or binocular if they had been monocular).

We assumed that every participant would have about the same experience that we had. We knew that they would be suspicious and that they might be reluctant to get things wrong, but we still expected some small effect to be evident in their judgments. As we collected more and more data, it was stunning to see that there was no obvious effect. Contrary to our expectation, some people gave slope estimates that were higher in the binocular condition than in the monocular condition. In most cases, there was very little difference between the two conditions in either the palm board data or in the verbal reports. We knew we were in a deep and disturbing kind of methodological vortex, however, when a student who had just given verbal estimates that differed by only 5 degrees in the two conditions, turned to one of us as he was leaving at the conclusion of the experiment and said (sincerely, it seemed): "It's a really big effect isn't it!" He had seen it too, but his numeric estimates certainly made this hard to know.

In the end we tested 26 participants, dividing them roughly evenly across which hill they saw first (one hill was 18.5 deg, the other was 21.5 deg) and whether they saw it binocularly or monocularly. Overall, if we look at both judgments of both slopes from each participant (i.e. attempt to measure the effect within-subjects), our average palm board matches were 31.9 deg (monocular) and 31.1 deg (binocular), and our average

verbal judgments were 45.7 deg and 42.8 deg, which were not reliably different from each other,  $t(26) < 1$ .

By looking only at the first trial each participant engaged in, however, we might hope to see the unprotected perceptual error as a between-subject effect. Indeed, first-trial palm board estimates (mean = 33.3 deg) in the monocular viewing condition were higher than first-trial palm board estimates in the binocular viewing condition (mean = 27.2 deg),  $t(25) = 2.03$ ,  $p = .027$ , one-tailed. The verbal estimates in the two conditions (50.0 deg monocular and 42.1 deg binocular) were not reliably different from each other,  $t(25) = 1.40$ ,  $p = .087$ , one-tailed, though they trended in the predicted direction.

We believe that even these between-subject comparisons understate the perceptual difference we observed. Thus even these judgments may reflect cognitive corrections for viewing state. The comparison of between-subject and within-subject differences suggests that in spite of our clear requests to these participants that they tell us about their perceptions, the apparent within-subjects constancy we measured across viewing states was due to strategic compensatory judgments. Our method, which used very similar slopes for both viewings, did nothing to prevent this. The idea that apparent constancy can come from cognitive corrections is well-supported (Granrud, 2009).

## **Experiment 2: Judgmental bias masquerading as perception**

On the other hand, sometimes we can find differences in judgments that may not be perceptual. An exciting new form of ecological theory of surface perception has been introduced by Proffitt and colleagues (Bhalla and Proffitt 1999; Proffitt 2006; Proffitt *et al.* 1995; 2003), suggesting that our perceptual experience has embedded in it aspects of our behavioral potential including our current physiological state. The new theory is exciting because it correctly notes that perceptions do not have to be geometrically accurate to be useful for planning actions. Moreover it supposes, consistent with Milner and Goodale (1995), that much of what is evident in conscious perception is there for longer-term planning (in minutes or hours) whereas the online control of precise action might be guided by unconscious visual information (in seconds or less). The theory is supported by a variety of interesting results and paradigms that, together, point to a conscious visual experience that is richly textured by the intentions, attitudes, emotions and energetics of the observer and the situation. From an evolutionary perspective, the intrusion of these sorts of information into one's conscious visual awareness can be motivated because (a) it is only conscious experience (rather than motor control) that is affected, and (b) these kinds of considerations might well be desired to be integrated into our perceptual experience for the purposes of efficient (effortless) planning and decision making.

A drawback of the theory is that it suggests that violations of constancy are not only ubiquitous, but also variable, depending on many factors. This could make perceptual prediction processes difficult. Moreover, the theory seems to blur the distinction between perception and judgment. For example, in an extension to cognitive dissonance, participants required to wear a ridiculous costume while walking in a public space judged the distance of the space to be larger if they were paid well for their efforts than if they were not (Balcetis and Dunning 2007). Classically, such biasing of judgments in the cognitive dissonance literature is not always regarded as perceptual, but the new theory encourages judgments of this sort to be classified as perceptions. Although the boundary is not always clear, we suggest that a distinction between perceptual effects and effects on judgment may be reasonably sustained in many relevant cases.

Indeed, we will present evidence that experimental demand characteristics can influence perceptual judgments without necessarily reaching down into perceptual experience. This is not to argue that there are no cases where perceptions may be affected by intentions, emotions or other extra-visual factors (e.g. Durgin and Gigone 2007; Durgin *et al.* 2005). It is likely, for example, that attentional factors can alter perceptual experience dramatically (Carrasco *et al.* 2004), and clearly physiological effects of age and fatigue, for example, may influence the quality of visual processing rather directly. Rather than doubting the possibility of perceptual effects, our immediate concern is with whether transient manipulations of such things as physical load (or embarrassment, or what have you) actually affect the perception of distance and geographical slope rather than merely judgments concerning these perceptions.

One of the most interesting (and ultimately controversial) effects reported by Proffitt *et al.* (2003) was that distances appeared greater when wearing a heavy backpack. A number of labs immediately began playing with this effect, including ours, and found it difficult to replicate (Hutchison and Loomis 2006; Woods *et al.* 2009). In the published controversies about this effect, several further claims became established which appear to contradict either the evolutionary account above or to contradict the equation of judgment and perception. Specifically, although within-subject designs are decidedly more sensitive in the face of inter-subject variability, Proffitt *et al.* (2006) have argued that Hutchison and Loomis (2006) failed to detect the backpack effect with a within-subject design because the scaling applied to the participants' perceptions in one condition would have carried over to the other. However, if participants can so easily undermine the alleged effect of the backpack, this would seem to undermine, in turn, the evolutionary value of the purported immediacy of these effects.

The claim that within-subject designs may introduce metacognitive contamination is not without merit on its own terms, as we have argued in the previous section, but it can only be supported insofar as a distinction is maintained between judgment and perception that is not maintained in the case of the cognitive dissonance results discussed above, for example. (Otherwise, we must conclude that backpacks really did not affect distance perception in the Hutchison and Loomis study, because they did not affect judgments.) This point aside, Proffitt *et al.* (2006) have correctly argued that whereas Hutchison and Loomis had failed to replicate the effect even with a between-subject design, the number of participants used in the experiment was less than that used by Proffitt *et al.* (2003), and the data appear to trend in the predicted direction. In other words, the failure to replicate in the between-subject version was inconclusive.

We set out to do the opposite of what Hutchison and Loomis had done. That is, rather than failing to replicate an experiment demonstrating a backpack effect on distance, we sought to experimentally produce a backpack effect. Our goal, however, was to test whether the effect of the backpack might be due to a judgmental bias in response to implicit demands of the experimental situation.

Demand characteristics of an experiment are the cues that participants receive as to what the experimental hypothesis might be (Orne 1959). The relationship between experimenter and participant is typically one in which participants perceive it to be their duty to help the experimenter by being cooperative. When participants receive cues as to how they are supposed to behave in an experiment, they often will behave in a manner consistent with this demand character (Orne 1962).

To test whether backpacks impose an experimental demand in a backpack experiment we first administered a brief survey: Thirty-one Swarthmore College undergraduates enrolled in Introductory Psychology were given the survey after completing an unrelated experiment. In the survey, an experiment similar to that of Proffitt *et al.* (2003) was described in which an experimenter has participants wear a

heavy backpack and make distance judgments (only the backpack condition was described). The question in the survey simply asked respondents to report what they thought the experimenter's hypothesis was in the experiment just described; no alternatives were presented. Twenty (65%) respondents indicated that the experimenter hypothesized that the backpack would affect distance judgments, and of those respondents, 16 (80%) described a hypothesis that involved distance judgments increasing as a result of wearing the backpack (the second most common hypothesis was simply that the backpack would in some way degrade performance). Thus, to the majority of our respondents, the experimental hypothesis actually entertained by Proffitt *et al.* (2003) was transparent, and it seems likely that for many of the participants in the original study, the hypothesis was similarly transparent.

The experiment we report here involved judgments of slope rather than distance, but the concerns are the same. Although effects of backpacks on perceptual judgments of slope have been reported (Bhalla and Proffitt 1999) and widely cited, we were surprised to learn that they had never been demonstrated in a controlled experiment. Instead, Bhalla and Proffitt measured slope perception in Introductory Psychology students who were all required to wear backpacks while making slope judgments. Bhalla and Proffitt then compared these judgments with previously published data they had collected with a different set of participants (passersby) in a different social context. Because their manipulation was not applied to equivalent groups, it was not a true experiment. Thus, the experiment we will describe here may be the first direct experimental test of the effect of backpacks on slope perception. We have subsequently replicated the result with a more sophisticated design by utilizing a post-experiment questionnaire (Durgin *et al.* 2009).

Crucially, we used three between-subject conditions in our experiment, rather than two, because it was essential to our design that we manipulate the presence or absence of an experimental demand as well as the presence or absence of a heavy backpack. Our participants were randomly assigned to conditions. In the baseline condition, participants made slope judgments without any backpack. In the standard backpack condition, participants made slope judgments while wearing a heavy backpack. In the critical control condition, participants wore the same heavy backpack while making slope judgments, but were first given a plausible explanation for the backpack that was intended to remove the experimental demand.

To create a plausible explanation for the requirement to wear a heavy backpack, the experiment was done in an immersive virtual reality and the backpack was described as containing equipment crucial for the head-mounted display (HMD) that the participants all wore. Indeed, in the control condition, the video processor for the HMD was carried in the backpack (along with several heavy weights), and the cables between the processor and the HMD were made to appear short so that wearing the backpack seemed necessary for wearing the HMD. To further provide participants with an alternative hypothesis about the purpose of the experiment we showed them simulated slopes composed of two different types of texture.

Whereas Bhalla and Proffitt (1999) asked participants their weight and set the backpack weight to be 1/6-1/5 of this, we did not want to call attention to weight in the low-demand condition, so we used a standard backpack weight of 25 lbs (11.3 kg) for all participants. In the previous semester, while conducting a pilot experiment, we had determined that this weight was at least 1/6 of the weight of 94% of the female participants in our participant pool. Our recruitment strategy involved inviting randomly selected females from the Introductory Psychology pool to participate for credit. The women recruited did not know that they were selected for gender. Thirty female

undergraduates students were divided evenly among the three conditions (two additional participants were excluded for failing to follow instructions).

Bhalla and Proffitt (1999) used only two slopes, but collected three types of measure for each (verbal estimates, visual estimates – using an adjustable 2D angle representation, and an unseen hand-manipulated palm board). They reported that the backpack affected verbal estimates for the lower (5 deg) hill, and visual estimates for both (5 deg and 31 deg), but did not affect palm board estimates for either. Proffitt *et al.* (1995) have shown that visual and verbal measures tend to measure the same things. We used a verbal measure as well as a haptic matching task (palm board).

There were three between-subject conditions. In the Baseline condition, no backpack was worn. In the Control condition, subjects wore a heavy backpack that was described as containing the video apparatus for the HMD (it did contain that 3.1 kg apparatus, but also contained an additional 8.2 kg of weights); it had long cables that ran into it from a computer and other short cables that ran out to the HMD. In the Standard condition, the backpack contained only weights (totaling 11.3 kg), and no explanation was given for why it had to be worn; the video processor was placed on a nearby surface, and several dumbbell style weights were visible on another nearby surface to emphasize that weights were being used.

The apparent experimental manipulation (to help with the deception) was that different textures were used to cover the slopes in the virtual environment. Four slopes were tested in the main experiment, from shallow (7 deg) to steep (28 deg) by steps of 7 deg, and each slope was shown using two different textures, one of which was the primary texture, having well-defined texels, and the other a more abstract “grassy” texture. Prior to these measurements, there were five practice trials used to camouflage the limited number of actual angles tested and to allow students to get used to the different textures. The first slope presented was always in the primary texture type and at the center of the range of slopes to be presented (i.e., 17.5 deg). There followed four additional practice trials that varied in texture and presented slopes that were both higher than 28 deg and lower than 7 degrees so as to render the experimental range a subset of the range seen. After 5 practice trials and 8, randomly ordered experimental trials, the four slopes were again presented in the primary texture (in random order) for the haptic response. Thus, participants completed a total of 17 trials (13 verbal and 4 haptic).

The stimuli were presented stereoscopically in an nVis HMD with a nominal 60 deg diagonal field of view (approximately 39 deg vertical and 49 deg horizontal). A HiBall optical head-tracker provided sub-mm precision at 120 Hz. The scene was viewed from eye-height in stereo (using the participant’s measured pupillary separation) rendered and displayed at 60 Hz with 1280 x 1024-pixel resolution using custom OpenGL software. The display was immersive and compensated for all head-movements, corrected to eye position. The total lag was less than 50 msec. The orientation of a rigid plastic palm board, mounted on a tripod at about chest level, was monitored by a second HiBall tracker. The palm board was placed higher than in Proffitt *et al.* (1995) because in this raised position it was easier to manipulate (He *et al.* 2007); the HMD ensured the palm board could not be seen.

Each virtual hill presented in the experiment was defined as a planar surface that extended above the line of sight of the observer and extended to the left and right farther than the observer could see. The hill surface smoothly curved over a meter of surface into the simulated ground surface on which the participants stood. To prevent inspection of the cross-section of the hill, observers viewed the hill through a virtual doorway that obstructed their view of the hill beyond approximately 54 degrees of azimuth to the left and right. The participant stood  $4.5 \pm 0.5$  m from the base of the hill, behind and

between two virtual walls. The walls were 8 m high, 0.25 m deep and were positioned 2 m in front and 2 m to the side of where participants stood. The height of the hill was always higher than eye-height and was varied so that the angle of gaze to the top of the hill did not vary consistently with hill slope.

Because the field of view inside the HMD was limited, participants were instructed during the first practice trial to look to the left and right before making their judgments to get a better sense of the spatial layout. They were also encouraged to look at the ground to their left and right to help stabilize their sense of what a horizontal surface looked like. Normally this information would be present in peripheral vision.

Finally, the participant was asked to provide a verbal estimate of the slope of the hill in degrees. This number was entered by the experimenter, and the virtual world went blank for about a second before the next hill was presented. After the 13 verbal trials were completed, the experimenter explained the use of the (unseen) palm board and had the participant reach out to it. The participants, who were encouraged to explore the hill visually as before, then adjusted the palm board and indicated when they felt that it was parallel with the slope of the hill. The final position was recorded. After four palm-board trials, (all with the primary texture on the surface), the HMD was removed and the experimenter fully debriefed the participants. The entire procedure took about 20 minutes.

A graph of the mean verbal estimates and palm board estimates for each of the three conditions is shown in Figure 1. The verbal means for the Baseline, Control, and Standard conditions were, respectively,  $29.3 \pm 15$  deg,  $28.5 \pm 5.2$  deg, and  $36.7 \pm 8.7$  deg. Consistent with the demand hypothesis, verbal estimates in the Standard condition were reliably greater than those in the Control condition,  $t(18) = 2.578$ ,  $p = .019$ . That is, when the backpack was worn as part of the experiment (without explanation), slopes were judged reliably steeper than when the same weighted backpack was worn (and described) as an incidental part of the apparatus. Consistent with the report of Bhalla and Proffitt (1999) there was no reliable effect of condition on the palm board settings (but see Durgin *et al.* 2010).

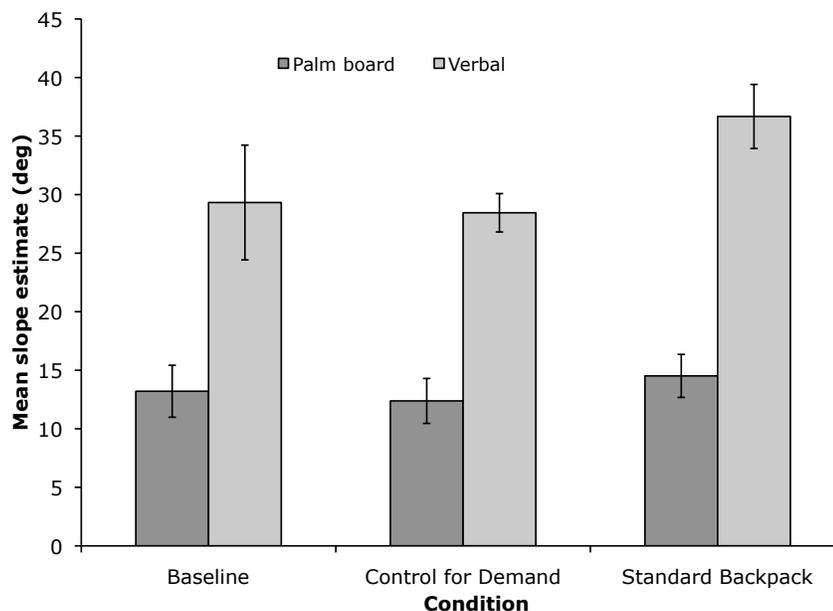


Figure 1. Mean slope estimates in Experiment 2, in which virtual slopes were judged by three groups of participants who either wore no backpack (Baseline),

wore a heavy backpack that was explained as part of the video processing apparatus for the head mounted display (Control for Demand) or wore a heavy backpack simply at the request of the experimenter while making the judgments (Standard Backpack).

So do heavy backpacks affect the perception of slope, or do they affect only the estimates of the participants? Having conducted a true experiment on the effects of wearing a backpack on slope, we have found that the experimental demand posed by the backpack in the Standard condition is sufficient to produce an effect on verbal judgments (though not on haptic slope matches). The effect is of the same magnitude reported by Bhalla and Proffitt (1999). We can conclude that the effect on verbal judgments is due to experimental demand (and thus most likely an effect on judgment, not perception) rather than the weight of the backpack, because we have a control condition in which the same backpack is worn, but a plausible explanation is provided for wearing it.

Based on our data, it seems reasonable to conclude that effects of backpacks on perceptual judgments are due to cognitive biases induced by the social context of the experiment rather than effort-based changes in perception. Notice that because we have replicated the result reported by Bhalla and Proffitt (1999) and shown that it depends on demand characteristics, the details of our implementation of the experiment are not really at issue. Unlike studies that have failed to replicate backpack effects on distance (Hutchison and Loomis 2006; Woods *et al.* 2009), our goal has not been to argue that there are no effects of backpacks, but to show that the effects that have been reported so far are probably due to social influences on judgment rather than physiological influences on perceptual experience. In subsequent investigations we have found that a “compliant” subset of participants drive the effect: they give high slope estimates, are able to articulate the hypothesis afterward, and also state that they believe they were affected (Durgin *et al.* 2009). If they had been affected perceptually, they should have no way of knowing they had been affected.

If perceptual experiences reflect perceptual prediction, as has been argued for self-motion perception (Durgin 2009), then perceptual experience had better not be arbitrarily plastic. Effort theorists have tended to sidestep this concern by appealing to Milner and Goodale's (1995) separation of vision for action and vision for perception (e.g. Proffitt 2006). They describe palm boards as action measures. But the claim that, for example, palm board measures are “action measures” (Bhalla and Proffitt 1999) has little to recommend it; adjusting the consciously-perceived haptic orientation of a palm board by hand has no evident relationship to the motor action of stepping onto a hill (see Durgin *et al.* 2010). We suggest that palm board measures are simply less affected by judgmental biases (though they may still suffer from them) – as was also evident in Experiment 1. Our data support the conservative view that perceptual experience is probably not as subject to fluctuation as the effort theorists have argued. That is, whereas perceptual judgments (especially from memory) are subject to vagaries of social expectation and cognitive dissonance, there seems to be little evidence that perceptual experience is affected by transitory burdens.

### **Conclusion: Downhill from here**

When is perceptual constancy important? Li and Durgin (2009) have recently observed a striking apparent failure of constancy in the perception of downhill slopes: For hills or even small ramps viewed from the top, perceived slope is much steeper if one stands back from the edge of the hill so that one's incident gaze is nearly parallel

with the sloped surface. This means that if one approaches a steep (e.g. 20 deg) downhill slope from the level ground above it, the hill surface initially appears particularly steep, but then grows visibly shallower as one nears the edge. The same effect can be observed when approaching a flight of stairs. What was most striking to us in making this discovery, however, was that we had to look for it, and that no one else seems to have reported it before.

Whereas others have argued that downhill slopes are judged steeper than uphill slopes (Proffitt *et al.* 1995), we would now argue that there is no unique value for perceived slope from the top of a hill. It depends where you stand. We discovered this because we wondered whether aiming one's gaze down along a hill would help one to see the true orientation of the hill. This led us to explore viewing positions that were different distances from the edge and to notice that the apparent slope of the hill seemed even steeper (and therefore less accurate) as we stood a few steps back. We have found that this can be quantitatively modeled by a combination of proprioceptive error regarding gaze direction and logarithmic coding of optical slant – surface slant relative to gaze orientation. We have also found that the proprioceptive perception of head pitch (even with closed eyes) is greatly exaggerated (Li and Durgin 2009).

How can such a failure of orientation constancy go unnoticed in daily life? Our argument here is that the experience of virtual perceptual constancy sometimes depends on the predictability of the perceptual consequences of our actions. Much as we seldom notice the optic flow of the environment as we move – it is expected – even large fluctuations in apparent surface orientation may be unremarkable to us. Future work can seek to determine whether, in this particular case, this is because these apparent deformations are predictable perceptual consequences of our actions (Durgin *et al.* 2007b), as seems to be the case for the optic flow of the ground plane (Durgin *et al.* 2005; Durgin 2009), or because our own self-motion often masks apparent object rotations anyway (Wallach *et al.* 1974).

Throughout this paper we have sought to support the notion that there are perceptual facts that are distinct from judgments we make about them and we have pointed to correspondences as well as discrepancies between introspective experience and measurable performance. As an example of a correspondence, we reviewed evidence that motion parallax does not seem to support the perception of empty space in the same way that binocular stereopsis does at intermediate distances. In Experiment 1 we suggested that real perceptual differences in surface slope were being masked by metacognitive strategies that produced null effects in within-subject comparisons. Between-subject comparisons revealed the predicted (and probably real) perceptual differences in slope perception for monocular viewing compared to binocular viewing, though likely underestimated them. These differences corresponded qualitatively with our own subjective impressions. With Experiment 2, however, we argued that judgmental biases rather than perceptual differences were responsible for the effects of backpacks on judgments of slope. When a heavy backpack was worn in a context that licensed the implication that the backpack was intended to affect slope judgments, slope estimates were higher than when the same heavy backpack was worn in a context that removed this demand character of the experimental context.

Whereas traditional constancy research has often confounded metacognitive judgments and perceptual experience, Granrud's (2009) work suggests that underconstancy is more the rule than the exception when metacognition is directly assessed. Here we have advanced the notion that stable underconstancy may be functional in supporting the guidance and control of action because it preserves the structured correlations (such as between self-motion and perceived optic flow) that can

be used to tune perception most precisely for the control of action (Durgin 2009; Durgin *et al.* 2010).

## References

- Abrams, Alicia B., James M. Hillis, and David H. Brainard (2007). The relation between color discrimination and color constancy: When is optimal adaptation task dependent? *Neural Computation* 19: 2610–37.
- Balçetis, Emily, and David Dunning (2007). Cognitive dissonance and the perception of natural environments. *Psychological Science* 18: 917–21.
- Beusmans, Jack M. H. (1998). Optic flow and the metric of the visual ground plane. *Vision Research* 38: 1153–70.
- Bhalla, Mukul, and Dennis R. Proffitt (1999). Visual-Motor recalibration in geographical slant perception. *Journal of Experimental Psychology: Human Perception & Performance* 25: 1076–96.
- Carrasco, Marisa, Sam Ling, and Sarah Read (2004) Attention alters appearance. *Nature Neuroscience* 7: 308–13.
- Durgin, Frank H. (2009). When walking makes perception better. *Current Directions in Psychological Science* 18: 43–7.
- Durgin, Frank H., Jodie A. Baird, Mark Greenburg, Robert Russell, Kevin Shaughnessy, and Scott Waymouth (2009). Who is being deceived? The experimental demands of wearing a backpack. *Psychonomic Bulletin & Review* 16: 964–9.
- Durgin, Frank H., Laurel Evans, Natalie Dunphy, Susan Klostermann, and Kristina Simmons (2007a). Rubber hands feel the touch of light. *Psychological Science* 18: 152–7.
- Durgin, Frank H., and Krista Gigone (2007). Enhanced optic flow speed discrimination while walking: Multisensory tuning of visual coding. *Perception* 36: 1465–75.
- Durgin, Frank H., Krista Gigone, and Rebecca Scott (2005). Perception of visual speed while moving. *Journal of Experimental Psychology: Human Perception and Performance* 31: 339–53.
- Durgin, Frank H., Alen Hajnal, Zhi Li, Natasha Tonge, and Anthony Stigliani (2010). Palm boards are not action measures: An alternative to the two-systems theory of geographical slant perception. *Acta Psychologica* 134: 182–97.
- Durgin, Frank H., Catherine Reed, and Cara Tigue (2007b). Step frequency and perceived self-motion. *ACM : Transactions in Applied Perception* 4/1.5: 1–23.
- Gibson, James J. (1950). *The perception of the visual world*. Cambridge: The Riverside Press.
- Gibson, James J., and Janet Cornsweet (1952). The perceived slant of visual surfaces—optical and geographical. *Journal of Experimental Psychology* 44: 11–15.
- Granrud, Carl E. (2009). development of size constancy in children: A test of the metacognitive theory. *Attention, Perception, & Psychophysics* 71: 644–54.
- He, Zijiang J., Ji Hong, and Teng Leng Ooi (2007). On judging surface slant using haptic (palm-board) and verbal-report task [Abstract]. *Journal of Vision* 7: 282a.
- Hillis, James M., Marc O. Ernst, Martin S. Banks, and Michael S. Landy (2002). Combining sensory information: Mandatory fusion within, but not between, senses. *Science* 298: 1627–30.
- Hutchinson, Jeffrey J., and Jack M. Loomis (2006). Does energy expenditure affect the perception of egocentric distance? A failure to replicate experiment 1 of Proffitt, Stefanucci, Banton, and Epstein (2003). *The Spanish Journal of Psychology* 9: 332–9.

- LeClair, Andrew, and Frank H. Durgin (2008). Depth interval perception: Comparing binocular stereopsis with motion parallax in "action space" [Abstract]. *Journal of Vision* 8: 857a.
- Li, Zhi, and Frank H. Durgin (2009). Downhill slopes look shallower from the edge. *Journal of Vision* 9/11.6: 1–15.
- Mamassian, Pascal (2008). Depth, but not surface orientation, from binocular disparities [Abstract]. *Journal of Vision*, 8, 89a.
- McGurk, Harry, and John MacDonald (1976). Hearing lips and seeing voices. *Nature* 264: 746–8.
- Milner, A. David, and Melvyn A. Goodale (1995). *The visual brain in action*. Oxford: Oxford University Press.
- Nachmias, Jacob (2006). The role of virtual standards in visual discrimination. *Vision Research* 46: 2456–64.
- Nawrot, Mark (2003). Eye movements provide the extra-retinal signal required for the perception of depth from motion parallax. *Vision Research* 43: 1553–62.
- Noë, Alva (2004). *Action in Perception*. Cambridge: MIT Press.
- Ono, Mika E., Josée Rivest, and Hiroshi Ono (1986). Depth perception as a function of motion parallax and absolute-distance information. *Journal of Experimental Psychology: Human Perception and Performance* 12: 331–7.
- Orne, Martin T. (1959). The nature of hypnosis: Artifact and essence. *Journal of Abnormal and Social Psychology* 58: 277–99.
- (1962). On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American Psychologist* 17: 776–83.
- Proffitt, Dennis R. (2006). Embodied perception and the economy of action. *Perspectives on Psychological Science* 1: 110–22.
- Proffitt, Dennis R., Mukul Bhalla, Richard Gossweiler, and Jonathan Midgett (1995). Perceiving geographical slant. *Psychonomic Bulletin & Review* 2: 409–28.
- Proffitt, Dennis R., Jeannine Stefanucci, Thomas Banton, and William Epstein (2003). The role of effort in perceiving distance. *Psychological Science* 14: 106–13.
- Proffitt, Dennis R., Jeannine Stefanucci, Thomas Banton, and William Epstein (2006). Reply to Hutchinson and Loomis. *The Spanish Journal of Psychology* 9: 340–2.
- Richards, Whitman (1985). Structure from stereo and motion. *Journal of the Optical Society of America* 2A: 343-9.
- Sacks, Oliver (2006). A neurologist's notebook: Stereo Sue. *The New Yorker*, June 19, 64–73.
- Wallach, Hans, Linda Stanton, and Dean Becker (1974). The compensation for movement-produced changes in object orientation. *Perception & Psychophysics* 15: 339–43.
- Wolfe, Jeremy M., Keith R. Kluender, Dennis M. Levi, Linda M. Bartoshuk, Rachel S. Herz, Roberta L. Klatsky, and Susan J. Lederman (2006). *Sensation & Perception*. Sunderland, MA: Sinauer.
- Woods, Adam. J., John W. Philbeck, and Jerome V. Danoff (2009). The various "perceptions" of distance: An alternative view of how effort affects distance judgments. *Journal of Experimental Psychology: Human Perception and Performance* 35: 1104–17.