Journal of Experimental Psychology:
Human Perception and Performance
2015, Vol. 41, No. 4, 929–939

0096-1523/15/$12.00 http://dx.doi.org/10.1037/a0038709

Stochastic or Systematic? Seemingly Random Perceptual Switching in Bistable Events Triggered by Transient Unconscious Cues

Emily J. Ward and Brian J. Scholl
Yale University

What we see is a function not only of incoming stimulation, but of unconscious inferences in visual processing. Among the most powerful demonstrations of this are bistable events, but what causes the percepts of such events to switch? Beyond voluntary effort and stochastic processing, we explore the ways in which ongoing dynamic percepts may switch as a function of the content of brief, unconscious, independent cues. We introduced transient disambiguating occlusion cues into the Spinning Dancer silhouette animation. The dancer is bistable in terms of depth and rotation direction, but many observers see extended rotation in the same direction, interrupted only rarely by involuntary switches. Observers failed to notice these occasional disambiguating cues, but their impact was strong and systematic: Cues typically led to seemingly stochastic perceptual switches shortly thereafter, especially when conflicting with the current percept. These results show how the content of incoming information determines and constrains online conscious perception—even when neither the content nor the brute existence of that information ever reaches awareness. Thus, just as phenomenological ease does not imply a corresponding lack of underlying effortful computation, phenomenological randomness should not be taken to imply a corresponding lack of underlying systematicity.

Keywords: bistable images, ambiguous figures, visual awareness, Spinning Dancer illusion, silhouette illusion

What we see is a function not only of incoming stimulation, but of unconscious inferences in visual processing (Helmholtz, 1866/1925). It has been argued that this is true of all forms of visual perception, since any pattern of light that strikes the eyes is always consistent with multiple possible local environments (e.g., Gregory, 1980; Rock, 1983). Perhaps the most powerful demonstrations of this, however, are multistable images, wherein the same stimulus alternates between (at least) two very different percepts, corresponding to two competing stable states of an underlying dynamical system. Almost any visual cue can fuel multistability in this way, for example, due to ambiguity in depth (as in the Necker cube), orientation (as in the duck/rabbit figure), figure/ground status (as in the Rubin face/vase), motion direction (as in apparent motion quartets), facing direction (as in Attneave triangles), and many other cues (for a review of these and many more cases, see Long & Toppino, 2004). In all such cases, at least part of the content of our percepts must stem from internal factors, since the stimuli themselves are unchanging. Accordingly, a great deal of research has explored the factors that influence perceptual switching in multistable images.

What Causes Percepts to Switch?

Perceptual switching in multistable images can be triggered by multiple factors. First, and perhaps most intuitively, such switching can seemingly occur for no reason at all—for example just due to noise in the shifting dynamics of network representations with multiple stable states. This possibility is supported by the phenomenology of seemingly haphazard and unintended switches, and by the evidence that periods of dominance and suppression are characterized by stochastic independence (Borsellino, De Marco, Allazetta, Rinesi, & Bartolini, 1972; Fox & Herrmann, 1967; Taylor & Aldridge, 1974; Walker, 1975).

Second, perceptual switching can be systematically influenced by many different intrinsic properties of the stimulus itself. Consider a structure-from-motion stimulus, for example, in which an array of moving dots can be seen as a transparent three-dimensional sphere, rotating to either the left or the right. If the strength of one interpretation is enhanced—say, by simply making the dots moving in one direction brighter than the others—then the matching percept will tend to dominate (e.g., Klink, van Ee, & van Wezel, 2008). Similarly, simply seeing one interpretation for an extended period will induce adaptation, making switches to the competing interpretation more likely, across multiple timescales (e.g., Harrison, Backus, & Jain, 2011; Jackson & Blake, 2010). Of course, these sorts of stimulus-driven factors are in no way inconsistent with the seeming stochastic nature of perceptual switching: These may

---

Emily J. Ward and Brian J. Scholl, Department of Psychology, Yale University.

For helpful conversation and comments on earlier drafts, we thank Gary Lupyan. For assistance with data collection, we thank Julie Qiu.

Correspondence concerning this article should be addressed to either Emily J. Ward or Brian J. Scholl, Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205. E-mail: emily.ward@yale.edu or brian.scholl@yale.edu

make certain types of switches more or less likely during particular intervals, but the exact nature and timing of any particular switch may still seem largely haphazard.[1]

Third, completely extrinsic visual stimulation can induce perceptual switching—perhaps "shocking" a dynamical system out of a stable state and allowing it to resettle into a new state. For example, simply presenting a brief (but conscious) visual transient can induce a subsequent perceptual switch while viewing phenomena such as a Necker cube or bistable apparent motion (Kanai, Moradi, Shimojo, & Verstraten, 2005).

Fourth, multistable images can sometimes switch due to voluntary effort—that is, because we explicitly try to make them switch. Observers can maintain a specified percept or can switch between percepts based on explicit instructions to do so (Toppino, 2003), and can willfully change their rate of switching (Meredith & Meredith, 1962). This influence of volition could occur relatively directly (e.g., by intentional priming of certain representations), or indirectly (e.g., via the voluntary differential allocation of attention to one part of a stimulus vs. another). As one example of the latter possibility, it seems that some intentional switching in the Necker cube is due to (possibly unconscious) attentional prioritization of (or even greater numbers of overt fixations on) specific contours and vertices (e.g., Peterson & Gibson, 1991a; Toppino, 2003).

Finally, switching can be influenced by explicit knowledge of the multistability itself, and of the possible competing percepts. Switching occurs more readily, for example, when the possible interpretations are explicitly recognizable (e.g., Peterson & Gibson, 1991b), and inversion of an otherwise-recognizable figure can dramatically reduce its perceptual dominance in bistable figure/ground relationships (Peterson & Gibson, 1994), binocular rivalry (Yu & Blake, 1992), and ambiguous line drawings (Girgus, Rock, & Egatz, 1977). At the extreme, in some cases there may be no perceptual switching at all when observers are unaware of the nature of one of the competing interpretations (Rock & Mitchener, 1992; but see Kosegarten & Kose, 2014; Mitroff, Sobel, & Gopnik, 2006).

Here we are interested in a possibility that does not fall cleanly into any of these particular categories, in which unconscious transient cues may lead to subsequent perceptual switches that nevertheless appear to observers to be completely haphazard. Unlike most studies of "top down" factors (e.g., voluntary effort or knowledge of ambiguity), this possibility involves specific bottom-up stimulus cues. Unlike studies that change the nature of the to-be-seen stimuli (e.g., Klink et al., 2008), the present experiments involve cues that are not themselves part of the resulting conscious percept. (It's not simply that observers fail to appreciate how and why the image cues are influencing their percepts; rather, they fail to see those image cues in the first place.) But unlike studies wherein conscious transients may prompt subsequent switching (e.g., Kanai et al., 2005), we are interested in the possibility that the specific nature of the cue may influence switching—so that, for example, Conflicting Cues may have a different influence than Consistent Cues, controlling for lower-level stimulus properties. Finally, whereas there are many other examples of unseen cues influencing what is seen later on, here we explore how an unseen transient cue may influence an ongoing percept, during extended online viewing. In sum, we aim to explore how the content of incoming information determines and constrains multistable perception—even when neither the content nor the brute existence of that information ever reaches awareness.

The Spinning Dancer

In principle, subtle transient cues could be used to disambiguate almost any multistable image. Such studies would be extremely difficult to realize in practice, however, given that the percepts are often so volatile and the switching so frequent. For example, in many common bistable images (such as the Necker cube or the Rubin vase), switches typically occur roughly every 5 s, but they can reach up to 1 switch/s when observers try to switch frequently (Windmann, Wehrmann, Calabrese, & Güntürkün, 2006). Moreover, even in dynamic stimuli wherein switching may be less frequent (such as certain structure-from-motion displays), such switches may only be likely during certain specific local stimulus configurations (Pastukhov, Vonau, & Braun, 2012). As a result, it would be difficult using any such stimuli to determine whether any particular switch was in fact driven by a bottom-up cue, or was just yet another frequent stochastic switch that coincidentally occurred shortly after the cue. (Note that this same challenge may not arise when the transient is allowed to be nonsubtle and fully conscious; cf. Kanai et al., 2005.)

A notable exception to the general rule of frequent switchability is the Spinning Dancer illusion: A spinning woman is depicted in silhouette, so that both her orientation in depth and her direction of rotation are bistable. Because this animation involves a silhouette, the depth ambiguity is apparent even in a single static frame, as in Figure 1: Here, you can see the dancer as either (a) facing you, with her left leg extended, or (b) facing away from you, with her right leg extended. The dynamic animation adds additional (though linked) ambiguity in rotation direction via the kinetic depth effect, but it also yields an unexpected side effect: Now the bistability seems somehow much more subtle, and the switching much less frequent. Indeed, it is striking that many observers who can otherwise make common ambiguous figures switch at will seem relatively powerless to do so with this animation—tending to see the dancer rotating in the same direction for long periods of time (despite continuous attempts to get her to switch), interrupted only rarely by involuntary switches (e.g., one recent paper reports that there are 5.18 switches/min with this stimulus, which is approximately half the rate of switching that occurs with most other bistable stimuli; Windmann et al., 2006).

This animation was developed in 2003 by the Japanese web designer Nobuyuki Kayahara (who initially called it the "silhouette

---

[1] Many studies of the role of stimulus factors in perceptual switching have involved binocular rivalry (BR)—for example, showing that the eye to which the "stronger" (e.g., brighter, or higher-contrast) stimulus is presented will tend to dominate (e.g., Levelt, 1966; Moreno-Bote, Shpiro, Rinzel, & Rubin, 2010). In general, perceptual switching during BR has many similar properties to switching in other multistable stimuli (Leopold & Logothetis, 1999), but we do not consider this phenomenon in the present paper given our focus on cases in which percepts can vary dramatically even while the input remains constant. In this context, BR seems like a special case: While the external stimulus may not be varying in most BR displays, selection during BR can occur as a result of interocular competition during very early in visual processing (e.g., Tong, Meng, & Blake, 2006), essentially changing the input to much of (downstream) vision—a situation that seems highly distinct from the stimuli employed here.

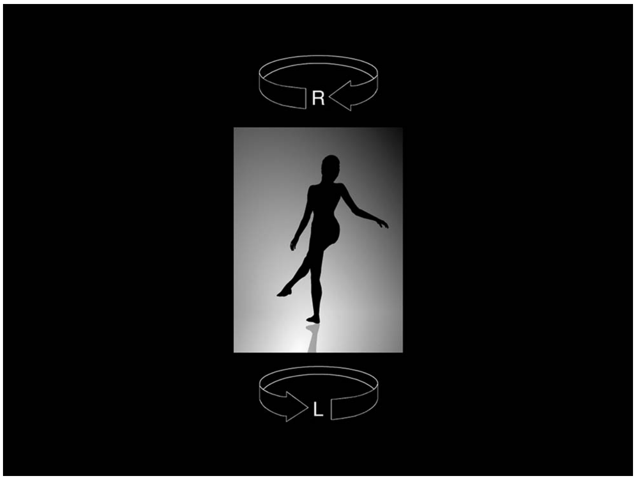

Figure 1. Observers monitored the Spinning Dancer animation and continuously indicated their current percept of her rotation direction (and thus also her 3-D pose) by holding down one of two keys.

illusion"[2]), and it has subsequently experienced some notoriety on the Internet, where it has been used for both explicitly nonscientific purposes (e.g., as an advertising lure) and for pseudoscientific purposes (e.g., as a purported test of "right-brained" vs. "left-brained" personalities; for discussion see Troje, in press). Here we refer to it as the Spinning Dancer illusion, though it has subsequently spawned many other variants.[3]

The seemingly unique properties of the Spinning Dancer display (perhaps along with its notoriety) have recently attracted the interest of several vision scientists, who have begun to explore the nature of this bistable animation. Most observers initially see the dancer as rotating clockwise (Troje, in press), and whatever initial direction is perceived tends to subsequently dominate (Liu, Tzeng, Hung, Tseng, & Juan, 2012). Some of these researchers have also pointed out that while the figure is a pure silhouette (and so has no depth cues involving self-occlusion of the dancer's limbs), there are more subtle depth cues available in the original animation. In particular, observers will tend to see ambiguous figures as if positioned below their viewpoint (the so-called viewing-from-above bias), and this combined in subtle ways with a visible shadow (and a slightly off-center viewing angle) may lead to the clockwise dominance in the original animation (Troje & McAdam, 2010; see also Troje, in press).

Despite these biases and the initial resistance to intentional switching, the Spinning Dancer display is bistable, and observers' percepts do switch. Although most of the infrequent switches seem to be purely stochastic, switching can be influenced by both bottom-up factors (e.g., with more switching at higher animation rates) and top-down influences (e.g., by intentionally focusing on the dancer's feet; Liu et al., 2012). Here, we exploit the infrequent switching of the Spinning Dancer display in order to ask about how subtle transient disambiguating image cues may influence perceptual switching.

The Current Study

A silhouette is ambiguous in depth precisely because it contains no internal depth cues such as self-occlusion relationships. Accordingly, a silhouette can be rendered formally unambiguous by providing such cues—for example, adding a contour to the Spinning Dancer to explicitly specify which leg is in front of the other, as depicted in Figure 2. (Because this stimulus is inherently dynamic, we also provide online animations at http://www.yale.edu/perception/dancer/.) Here we explore the influence of such explicit occlusion cues on perceptual switching, while (a) presenting the cues in a subtle and transient fashion so that observers are unable to report their existence, and (b) varying the specific nature of the cues to be either consistent or conflicting with the observer's current percept during extended viewing. If such unconscious cues influence the resulting percepts in a way that is contingent on their particular content, this would indicate that "bottom-up" visual cues can influence perceptual switching even when the cues do not figure into the resulting percepts, and, even while the switches feel to observers to be entirely stochastic.

General Method: Apparatus and Stimuli

Stimuli in Experiment 1 were presented online via Amazon Mechanical Turk, and so viewing distance was not controlled. Stimuli in Experiments 2–4 were presented at the center of a display subtending 34.71° × 26.72°, with observers situated (without head restraint) approximately 50 cm away (with all extents below calculated based on this viewing distance). Animations in Experiment 1 were presented using custom software written in jQuery. Animations in Experiments 2–4 were presented using custom software written in Python using the PsychoPy libraries (Peirce, 2007).

Given the enduring interest in the initial Spinning Dancer animation, we chose to use the original version of the illusion, which we obtained from Wikipedia (http://en.wikipedia.org/wiki/File:Spinning_Dancer.gif). This animation (8.85° × 10.00°) depicts a female figure as a black silhouette on a grayscale gradient background (with a visible shadow from the lower foot), viewed from a camera elevation of 6.8° above or below the horizontal plane (Troje & McAdam, 2010), with a maximal vertical and horizontal extent of 7.15°, rotating at 176.47°/s in the frontal depth plane, such that a full 360° rotation required approximately 2 s (2,040 ms—with each of 34 frames being presented for 60 ms, except in Experiment 1 and 3, as described below). Because of the lack of depth cues, the dancer can be perceived as rotating either clockwise or counterclockwise.

In Experiments 2–4, we presented this animation at the center of a black display with two 6.08° × 3.32° circular arrows (and the letters "R" and "L") presented, respectively, above and below the figure (as in Figure 1) in order to depict the two possible rotation directions (used, as described below, as response-mapping reminders).

Explicit contours were added at certain places and moments of the animation in order to disambiguate the orientation in depth (and the rotation direction) of the figure via the specification of self-occlusion cues. Examples of such contours are depicted in Figure 2, where the addition of a single white contour internal to

---

[2] As of this writing, the original animation remains available on the designer's website at http://www.procreo.jp/labo/labo13.

[3] As of this writing, for example, a feline variant (with a rather dramatic soundtrack) can be viewed at http://www.youtube.com/watch?v=KtSHsMTFfVg.

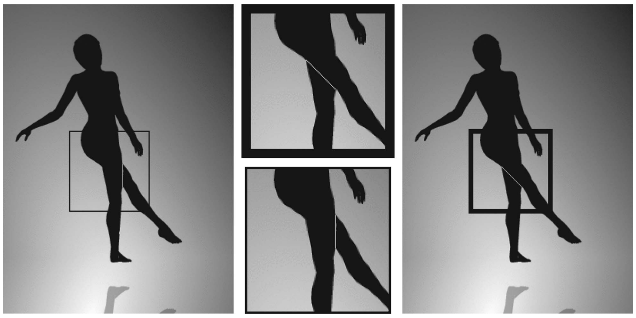

Figure 2. Brief disambiguating contours were flashed onto the Spinning Dancer animation to indicate occlusion relationships in what was otherwise a pure silhouette. In the same silhouette, for example, the contours could indicate either that the dancer was facing away from the observer (as in the leftmost panel) or toward the observer (as in the rightmost panel; with the central panels illustrating the cues themselves in close-up)—thus providing either conflicting or consistent information relative to the observer's current percept.

the silhouette disambiguates which of the dancer's legs is extended, and so determines the direction she is facing. These cues were taken directly from the Wikipedia entry for the illusion, where they were used (with no attribution) as a simple way of disambiguating the image.[4] For our purposes, this sort of disambiguation cue is perfect, since these contours (unlike fully textured disambiguations, as in Troje & McAdam, 2010), can be readily added and removed in a relatively subtle manner (as verified in Experiment 1).

As is apparent from the online demonstrations, all contours indicated occlusion relationships involving either (a) the dancer's outstretched leg (as intersecting either the thigh, knee, or calf of the other leg, depending on the rotation angle), or (b) one of the dancer's arms (as intersecting her torso). All contours were presented with a stroke of a single pixel, and with an extent that varied from 0.44° to 1.13°. Of the 34 constituent frames of the basic animation, only a subset could include a disambiguating contour in principle (14 frames for clockwise rotations; 21 frames for counterclockwise rotations), and on any given frame, only a single contour was ever visible at once (depending on the rotation angle), with the exception of a single frame where one contour of each of the two types described above was present. Frames with disambiguating contours were substituted for the equivalent ambiguous frames at certain moments during the animations as described below.

Experiment 1: Assessing Awareness

We aimed to test the influence of transient cues that were subtle enough that many observers would not consciously notice them, and would be unable to report them. Of course, it was certainly *possible* to notice cues such as those depicted in Figure 2, if observers were explicitly directed to them. But they seemed subtle enough that they could go unnoticed without such explicit direction. In this sense, the transient cues in the present studies served as types of "unexpected events," as in studies of inattentional blindness (e.g., Mack & Rock, 1998; Most, Scholl, Clifford, & Simons, 2005)—notice*able*, but nevertheless often "missed" and unreportable in practice.

In an initial independent test of how likely observers were to notice and report the white contour lines presented on the Spinning Dancer, observers (tested online) were told to simply view the Spinning Dancer animation and then to answer some questions about what they had seen. Observers viewed the animation in its fully ambiguous state, followed by the presentation of the disambiguating contours. These contours were presented as subtle white lines (as in Figure 2, and as employed in Experiments 2–4), but we also tested versions where the contours were colored bright red, in order to ensure that they could be reported when more salient.

Participants

Observers were recruited and run online via Amazon Mechanical Turk. (For discussion of this pool's reliability, see Crump, McDonnell, & Gureckis, 2013). Sixty observers were recruited—30 for each of the two conditions described below.

Stimuli

The display was identical to that described in the general method except as follows. To maximize the visibility of the contour cues, each frame of the animation was presented for approximately 85 ms, and each of the 14 possible contours frames was always presented during each rotation wherein they were present. (This value is approximate only because different observers were surely using different display hardware, whose refresh rates would have differentially constrained the specific durations for which frames could be shown; for discussion see Elze, 2010). The specific contours were always consistent with the clockwise-spinning percept (which is the one that tends to initially dominate; Troje, in press). The red contours were identical in all ways to their white counterparts (as described in the General Method), except for their color.

Procedure

Observers were instructed to watch the Spinning Dancer animation carefully because they would be required to answer questions about the animation when it ended. All observers saw 12 rotations (34.68 s) of the dancer. Two groups were tested. Observers in the red-only group first viewed the entirely (contour-free) ambiguous animation for six rotations (17.34 s), followed by six rotations containing the red contours. Observers in the white-then-red group first viewed the ambiguous animation for four rotations (11.56 s), followed by four rotations containing the white contours, and then four rotations containing the red contours. Thus, the observers in this group viewed the white contours for a total of 4.76 s (14 contour frames × 85 ms each × four rotations).

After the animation finished, observers were asked a series of seven questions to probe their awareness of the contour lines: (1) First, they were asked to describe what they had seen in a few sentences. (2) They were asked if they had seen any red lines during the animation—and if so to describe where they had appeared. (3) They were shown a still frame containing one of the red lines (similar to Figure 2), and were asked if they had noticed it.

---

[4] As retrieved on May 5, 2014, from http://en.wikipedia.org/wiki/Spinning_Dancer.

(4) They were asked if they had seen any other similar (but not red) lines—and if so to describe them. (5) They were asked explicitly if they had noticed a *white* line of this type during the animation—and if so to describe where they had appeared. (6) They were shown a still frame containing one of the white lines (similar to Figure 2), and were asked if they had noticed it. (7) Finally, they were given the option of guessing explicitly which test group they had been in: the group that saw only the red contours, or the group that saw white contours followed by red contours.

Results and Discussion

To be maximally liberal, we counted observers as having noticed the red contours if they responded positively to question 3, or if they mentioned red contours at all in response to either questions 1 or 2. Similarly, we counted observers as having noticed the white contours if they responded positively to question 6, or if they mentioned white contours at all in response to questions 1, 4, or 5. Only nine observers (30%) noticed white contour lines when they were present (in the white-then-red group), and no observers reported seeing white contour lines when they were not present (in the red-only group; $\chi^2(1) = 10.59$, $p = .001$, $\varphi = 0.42$), indicating that observers were not likely to mistakenly say that they had seen contours that were not in fact present. Despite this, there was no difference between the two groups in their ability to guess which group they had been in, via question 7, $\chi^2(1) = 1.83$, $p = .176$, $\varphi = 0.17$.

Observers noticed the red contours at a much higher rate—21 (70%) in the white-then-red group (a significantly greater percentage than noticed the white contours, $\chi^2(1) = 9.60$, $p = .002$, $\varphi = 0.40$), and 23 (76.6%) in the red-only group. This indicates that observers were perfectly able to answer positively to these sorts of questions when they did see the asked-about stimuli, thus strengthening the inference that negative answers truly reflected a failure to notice the asked-about stimuli. Even when limiting the pool of observers in the white-then-red group to those who did notice the red contours, only 7/21 (33.3%) noticed the white contours.

These results suggest that most observers fail to notice the subtle white contours, even with nearly a full 5 s of cumulative exposure—and even when we know (from the results of the red-only group) that observers will readily report such stimuli when they are seen. (This rules out the possibility, for example, that observers did see the subtle white contours, but just didn't appreciate that these were the sorts of things that were being asked about.) Although observers certainly could be induced to see such cues, apparently they are functionally blind to them in the present context.

General Method: Experiments 2–4

Participants

A different group of 20 college students (mean age = 20 years) with self-reported normal or corrected-to-normal acuity participated in Experiments 2 and 3 in exchange for course credit or a small monetary payment. In the absence of any previous studies that had used the kinds of manipulations developed for the present experiments, we began with the heuristic assumption that the required sample size would be roughly equivalent to those used in other studies of the Spinning Dancer illusion (e.g., Liu et al., 2012; Troje & McAdam, 2010)—and after obtaining positive results in Experiment 2, we retained the same sample size for Experiment 3 in order to promote comparisons across experiments. Experiment 4 added a new variable that cut the effective power in half, and so we doubled the sample size to 40. Because we were interested in the effect of unconscious cues on perceptual switching, we adopted an especially conservative criterion, and only analyzed data from observers who (via specific debriefing questions, described below) reported not having been aware of the existence of the transient occlusion cues at any point during the experiment. Obtaining 20 nonnoticers per experiment in Experiments 2 and 3 required running 35 observers for both. Obtaining 40 nonnoticers in Experiment 4 required running 54 observers (64.52% nonnoticing rate across the three experiments). Note that these nonnoticing rates were very similar to those obtained in Experiment 1, which assessed awareness via both (a) more elaborate questioning and (b) using additional questions (about the red contours) that confirmed the ability of such questions to tap noticing.

Procedure

Observers were first exposed to the ambiguous Spinning Dancer animation, and indicated which direction of rotation they perceived. The experimenter then showed observers unambiguous versions of the clockwise and counterclockwise rotations in sequence—presented via two dots indicating eyes. (This eye cue indicated which way the dancer's head was facing at any given moment, which thus disambiguated the animation as a whole. This step was necessary since some observers were unable to notice the ambiguity in the initial presentation.)

Observers were then instructed that during the experiment they were to continuously monitor the animation and to indicate at every moment using keypresses which direction they saw the dancer rotating (either clockwise, in which case they held down one key; or counterclockwise, in which case they held down a different key). Observers were told explicitly that some people tend to see the dancer switch frequently, while others almost never see any switching. They were reassured that either experience was acceptable, and that they should simply carefully attend to the dancer at all times, to not assume that the switching behavior would stay constant, and to simply report at each moment what they saw.

Observers viewed exactly 800 rotations in a single experimental session lasting approximately 45 min, with self-paced breaks after every 50 rotations. Thus, from an observer's perspective, the experiment consisted of one long task, with no explicit trials, per se. Nevertheless, we analyzed their responses as a function of certain discrete events for each experiment, such as when disambiguating contours of certain types were presented in the animation.

When disambiguating contours were presented, they were always displayed for exactly one 60-ms "cue frame" at a time (though this parameter was also varied in Experiment 3), surrounded by ambiguous frames. Note that this duration for which the white contour was visible is nearly 80 times shorter than the cumulative period of visibility in Experiment 1.

The timing of the cue frames was constrained by what observers were currently perceiving, such that there had to be at least one full rotation since the last change in what key observers were holding down before a cue frame could be presented. This meant that the specific frames selected as cue frames, their timing, and even their number throughout the study differed for different observers, as follows.

Each observer saw a total of 27,200 frames (800 rotations × 34 frames/rotation), from which 200 frames were randomly preselected (differently for each observer) as triggering frames—with 100 selected from the first half of the experiment (always excluding the first three rotations), and 100 from the last half (always excluding the final three rotations). Whenever a triggering frame was reached, the next cue frame was selected randomly from among all of the potential frames in the next full rotation from that point. Once the next cue frame was selected in this way, it was presented as scheduled unless either (a) the observer indicated that their percept had changed in the interim, or (b) the particular cue frame chosen would occur before one full rotation after the last keypress (where that keypress could still have occurred before the triggering frame); in either of these cases, the cue frame was abandoned, and the process was reset (without replacement) by waiting for the next triggering frame.

After their session, each observer completed a funneled debriefing procedure during which they were asked about their experiences and about any particular strategies that they had employed. To determine their awareness of the disambiguating contours, observers were asked, “Did you notice a white contour line presented on the dancer at any point?” Only observers who answered this question negatively were included in the analyses, and the resulting percentages of “unaware” observers could then be compared with the results of Experiment 1.

Experiment 2: Transient Unconscious Cues

To explore how switching might be influenced by transient unconscious cues, observers viewed the ambiguous Spinning Dancer animation onto which occasional disambiguating contours were presented for a single frame each. These contours were always chosen so as to conflict with the observer’s current percept (as specified by their ongoing keypresses). For example, if the observer currently saw the dancer’s extended leg facing toward them, then they might see a cue such as that in the leftmost frame of Figure 2—whereas if they currently saw the dancer’s extended leg facing away from them, then they might see a cue such as that in the rightmost frame of Figure 2. We were thus interested in whether the presentation of these cues would cause perceptual switches, relative to baseline switching that occurred independent of the cues.

Method

Observers viewed the Spinning Dancer animation with occasional conflicting contour cues presented as described in the general method. On average, observers saw 153.65 cue frames (*SD* = 16.62), separated by an average of 4.99 rotations (*SD* = 0.60) or 11.68 s (*SD* = 1.56).

Results and Discussion

Observers experienced an average of 122.90 switches throughout the course of the experiment (*SD* = 95.35), with an average latency between switches of 24.22 s (*SD* = 16.62). (One observer did not switch at any point during the experiment, however, and so they were excluded from all analyses that included latency as a factor.) Considering each quartile of the experiment as an artificial block, the number of switches did not increase as a function of block, $F(1, 19) = 3.44$, $p = .079$, $\eta_p^2 = 0.15$, but the average interswitch latency decreased as a function of block, $F(1, 18) = 10.512$, $p = .005$, $\eta_p^2 = 0.36$; $M_{block1} = 52.05 \pm 43.61$, $M_{block2} = 25.09 \pm 16.52$, $M_{block3} = 23.88 \pm 20.03$, $M_{block4} = 19.77 \pm 17.44$.

If a transient cue is indeed causing perceptual switching, we might naturally expect switching to be roughly temporally locked to the cue—that is, for the cue to lead to a perceptual switch shortly thereafter. To test this, we first measured the number of switches that occurred during the 8 s after each cue frame, breaking the 8-s response window down into 1-s bins. This served as a measure cue-driven perceptual switching. However, to determine whether the unconscious disambiguating cues led to perceptual switching, we had to compare the frequency and likelihood of switching when a cue had been presented to the switching that did or did not occur at all other moments when a cue had *not* been presented. Therefore, we also measured the number of switches that occurred during every other 8-s window that did not follow a cue frame. We used a moving window approach, in which the 8-s window was shifted by 1 s in time, starting at the beginning of the experiment through the end of the experiment. Every 8-s window was included unless it began with a cue frame (in which case it was counted as a cue-driven switch, as described above), but windows *including* a cue frame (and potentially any resultant switching) were included. Thus, this approach provided a within-subject baseline that includes all switching, both stochastic and cue-driven.[5]

These data are depicted in Figure 3, presented as a proportion of switching relative to the baseline switching (so that, e.g., a value of 2 indicates that there was twice as much switching relative to baseline). There is a clear declining trend in this figure (verified by linear regression: $\beta = -0.207$, $F(1, 19) = 88.69$, $p < .001$, $\eta_p^2 = 0.36$), such that most of the switching caused by the cue occurred in the first couple of seconds that followed. Indeed, the amount of switching above baseline was statistically reliable only for the first two bins, $t_{1second}(19) = 6.10$, $p < .001$, $d = 1.36$; $t_{2seconds}(19) = 3.15$, $p = .005$, $d = 0.70$. Unexpectedly, it appears that this increase in switching as a main effect of the presence of a cue was balanced out by a corresponding *decrease* in switching in the later parts of the 8-s window—such that there was actually *less* switching then, compared with baseline: At bin three, the amount of switching was no greater than baseline ($p = .52$), but for most of remaining bins, the amount of switching was less than baseline (p

---

[5] We also performed these analyses and obtained similar results using the frequency and timing of perceptual switching in an independent baseline group of subjects, who observed the Spinning Dancer animation in its initial (fully ambiguous, silhouette) form, without any disambiguating cues. However, because the amount of perceptual switching varies dramatically from person to person, the within-subject baseline as described here provides a better, if more conservative, comparison.

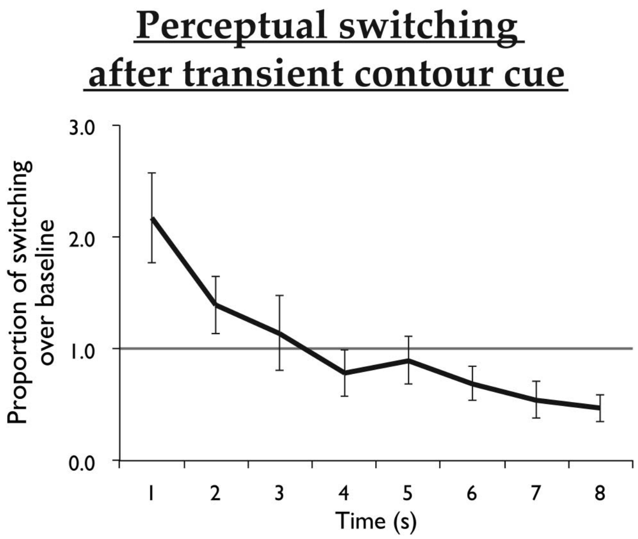

Figure 3. Observers' perceptual switching behavior in Experiment 2. Cue-driven switching as a function of the latency since the cue, presented as a proportion of switching above baseline (so that, e.g., a value of 2 indicates that there was twice as much switching relative to baseline). Cues led to more switching shortly afterward. Error bars indicate 95% CIs.

values < 0.01, except bin five which was not significant). (We discuss this result further in the general discussion.[6])

These analyses can be summarized in terms of the prominent trend depicted directly in Figure 3: Transient unconscious cues led to more perceptual switches (beyond a no-cue baseline) immediately afterward—for example with more than twice as many switches in the second following a cue.

Experiment 3: Cue Duration

One of the most striking aspects of Experiment 2 was that the disambiguating cues had such a robust influence on switching despite being so short. This led us to wonder whether even shorter cues would have a similar effect. The disambiguating contours in Experiment 2 were each presented for a single 60-ms cue frame, but our display hardware actually allowed for faster physical frame rates, such that single cue frames in Experiment 2 were actually presented for 6 screen refreshes (on a display refreshing at 100 Hz). Here we replicated Experiment 2 while testing a range of cue durations, from slightly longer durations than that used in Experiment 2, down to less than 17 ms.

Method

This experiment was identical to Experiment 2 except as follows. We tested 4 cue frame durations: 16.67 ms (i.e., 1 display refresh at 60 Hz), 33.33 ms (2 refreshes), 50.00 ms (3 refreshes), or 66.67 ms (4 refreshes). Each time a cue frame was selected (as described in the general method), one of these durations was chosen randomly. Thus each duration was presented approximately 35 times ($M_{16.67} = 35.10 \pm 8.20$, $M_{33.33} = 34.35 \pm 7.62$, $M_{50.00} = 34.35 \pm 8.80$, $M_{66.67} = 34.30 \pm 9.31$), with no significant differences in presentation frequency, $F(3, 57) = 0.29$, $p = .832$, $\eta_p^2 = 0.02$. The average separation between any two cues was 6.17 ± 3.89 rotations (15.95 ± 9.58 s).

Results and Discussion

Observers experienced an average of 203.20 switches throughout the course of the experiment ($SD = 162.09$), with an average latency between switches of 17.42 s ($SD = 10.47$). The number of switches increased as a function of block, $F(1, 19) = 18.88$, $p < .001$, $\eta_p^2 = 0.50$; $M_{block1} = 33.70 \pm 35.81$, $M_{block2} = 44.15 \pm 39.73$, $M_{block3} = 61.10 \pm 48.63$, $M_{block4} = 64.25 \pm 46.96$, and the average interswitch latency decreased as a function of block, $F(1, 19) = 28.83$, $p < .001$, $\eta_p^2 = 0.60$; $M_{block1} = 30.31 \pm 13.93$ s, $M_{block2} = 24.26 \pm 21.17$ s, $M_{block3} = 14.62 \pm 10.16$ s, $M_{block4} = 14.41 \pm 12.20$ s.

The number of cue-driven switches (those occurring within 2 s after the cue frame) is presented in Figure 4 broken down by cue duration, and presented as a proportion of switching above baseline (as computed by the same moving window analyses described in Experiment 2). The surprising result of this experiment is immediately clear from this figure: The duration of the unconscious contour line did not have any effect on the amount of subsequent switching, $F(3, 57) = 1.66$, $p = .187$, $\eta_p^2 = 0.08$. Although contour durations of 33.33 ms, 50.00 ms, and 66.67 ms produced switching significantly or marginally significantly above baseline, contour durations of 16.67 ms did not, $t(39) = 0.070$, $p = .945$, $d = 0.01$. This indicates that all but extremely short (16.67 ms) unconscious disambiguating cues are sufficiently powerful to cause observers to switch their percept of a bistable event.[7]

Experiment 4: Consistent Versus Conflicting Cues

So far, we have been assuming that the reason why a transient cue caused observers to switch what they saw is simply that cue always conflicted with the observer's current percept: If you saw the dancer facing away from you, then the momentary occlusion cue indicated that she was in fact facing toward you (thus necessitating a switch to bring the momentarily unambiguous input and the percept into alignment)—and vice versa. But of course, it is possible that the mere existence of the transient cue itself (beyond its particular content) could have led to increased switching (see Kanai et al., 2005)—perhaps by capturing attention in some way. We typically think of attention capture as an extremely conscious phenomenon, given the often tight link between attention and awareness. Nevertheless, it has been shown that attention can be robustly captured by certain transient cues even when observers are completely unaware of the existence and nature of those cues (e.g., Jiang, Costello, Fang, Huang, & He, 2006; Lin, Murray, & Boynton, 2009; Norman, Heywood, & Kentridge, 2013). As a result, in this experiment we compared the influence of transient occlusion cues that *conflicted* with the observer's current percept (as used in the previous experiments) with those that were *consistent* with the current percept (tested here in the same observers). If the influence of a transient unconscious cue depends on its specific

---

[6] The same pattern of results was obtained when analyzing the 15 observers who did notice the white contours, with no notable differences.

[7] The same pattern of results was obtained when analyzing the 15 observers who did notice the white contours, with no notable differences.

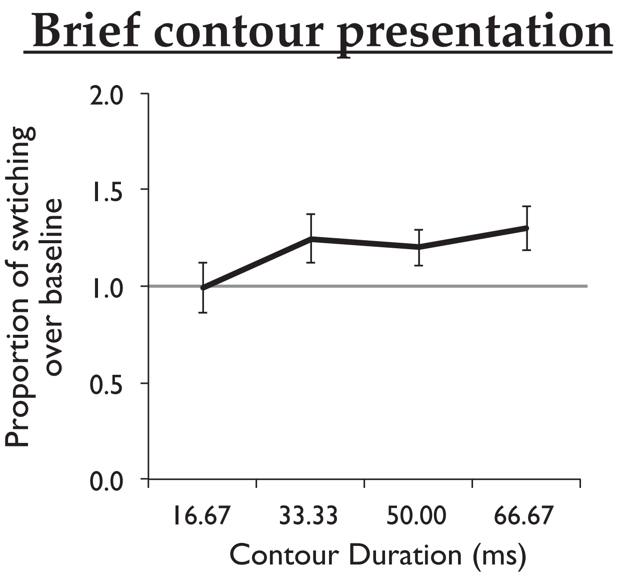

Figure 4. Observers' perceptual switching behavior in Experiment 3, showing roughly equivalent switching for all tested cue durations. All values are presented as proportions of switching above the baseline (so that, e.g., a value of 2 indicates that there was twice as much switching relative to baseline). Error bars are ± SEM.

nature, then we might expect more changes following Conflicting Cues compared with Consistent Cues.

Method

This experiment was identical to Experiment 2 except as follows. Half of triggering frames were randomly assigned to be potential Conflicting Cue frames, with the other half assigned to be potential Consistent Cue frames. (As noted in the general method, this effectively cut our number of trials per cell in half, and we compensated for this by running twice as many observers as in the previous experiments.) Because not all triggering frames were used as cue frames (as described in the general method), the total number of resulting cues varied slightly across observers, with approximately 70 cue frames per condition ($M_{consistent} = 70.35 \pm 13.50$; $M_{conflicting} = 70.75 \pm 11.52$), with no difference between the two, $t(39) = 0.42$, $p = .676$, $d = 0.07$. The average separation between any two cues was 5.25 ± 0.68 rotations (12.94 ± 3.26 s), and this separation did not vary by condition, $t(39) = 0.31$, $p = .755$, $d = 0.05$.

Results and Discussion

Observers experienced an average of 129.93 switches throughout the course of the experiment ($SD = 117.39$), with an average latency between switches of 27.39 s ($SD = 19.09$). The number of switches increased as a function of block, $F(1, 39) = 5.37$, $p = .026$, $\eta_p^2 = 0.12$; $M_{block1} = 24.80 \pm 30.52$, $M_{block2} = 26.75 \pm 23.08$, $M_{block3} = 34.00 \pm 40.19$, $M_{block4} = 44.38 \pm 51.80$, and the average interswitch latency decreased as a function of block, $F(1, 39) = 15.78$, $p < .001$, $\eta_p^2 = 0.29$; $M_{block1} = 52.00 \pm 42.85$ s, $M_{block2} = 31.44 \pm 24.40$, $M_{block3} = 28.53 \pm 23.94$, $M_{block4} = 25.13 \pm 21.78$.

We computed the degree of switching above baseline in the 8-s window following both Conflicting and Consistent Cues, and we again broke those responses down into 1-s bins. These data are depicted in Figure 5, again presented as a proportion of switching above baseline.

Inspection of this figure suggests three salient results. First, it appears that transient cues again led to a greater number of switches, especially in the window immediately following the cues (thus replicating the primary results of Experiment 2). Second, it appears that this increase in switching as a main effect of the presence of a cue (regardless of its content) was again balanced out by a corresponding decrease in switching in the later parts of the 8-s window. Third, and most importantly, we can see from the initial separation of the two lines that there was a general effect of Consistent Cues, but that there was an even greater effect of Conflicting Cues (which led to an especially stark increase in switching following the cues, up to nearly 50% more switching compared with the Consistent Cues for some of the early bins).

These impressions were all verified via the following statistical analyses. We first performed a two-way analysis of variance (ANOVA) with time and cue condition (consistent vs. conflicting) as independent variables on the normalized (i.e., above-baseline) switching frequencies. This yielded a main effect of time, $F(1, 39) = 114.10$, $p < .001$, $\eta_p^2 = 0.75$, but no main effect of cue condition, $F(1, 39) = 2.47$, $p = .124$, $\eta_p^2 = 0.06$. However, as observed in Figure 5, there was a significant Time × Cue Condition interaction, $F(1, 39) = 9.97$, $p = .003$, $\eta_p^2 = 0.20$. The main effect of time now included switching above baseline in the first bin, $t(39) = 5.69$, $p < .001$, $d = 0.90$, no switching beyond baseline for the second and third bins (both p values > 0.09), and the inhibitory component: Switching was significantly *below* baseline for the final five bins (all p values ≤ 0.001).

Most critically, we observed that Conflicting Cues led to more above-baseline switching than Consistent Cues for the first two bins (with a marginal effect in the 1st bin, $t_{1second}(39) = 1.91$, $p = .064$, $d = 0.30$; $t_{2seconds}(39) = 3.06$, $p = .004$, $d = 0.48$), but not for any of the other bins (all other p values > 0.49). Both Conflicting and Consistent Cues led to above-baseline switching in the first bin, but in second bin, only the Conflicting Cues led to above-baseline switching, $t_{1second}(39) = 2.78$, $p = .008$, $d = 0.44$. Consistent Cues did not lead to above-baseline switching in this bin, $t_{1second}(39) = 0.72$, $p = .477$, $d = 0.11$, or any of the others (which were numerically below baseline).[8]

The key result from this experiment can be summarized by noting that perceptual switching driven by unconscious transient cues depended here on the specific content of the cue (rather than by its mere presence).

Why was there more switching at all following Consistent Cues? We cannot be sure, but it seems plausible that this effect is simply due to attention capture effects driven by the transient itself. Indeed, extrinsic transients have been shown in past research

---

[8] When these analyses were run on the observers who noticed the white contour lines, the presence of a cue led to switching in the first bins, $t(11) = 3.19$, $p = .009$, $d = 0.92$, but there was no difference between Conflicting and Consistent Cues on switching, $t(11) = 0.80$, $p = .441$, $d = 0.23$. (Exactly 2 out of 14 subjects were excluded due to missing data.)

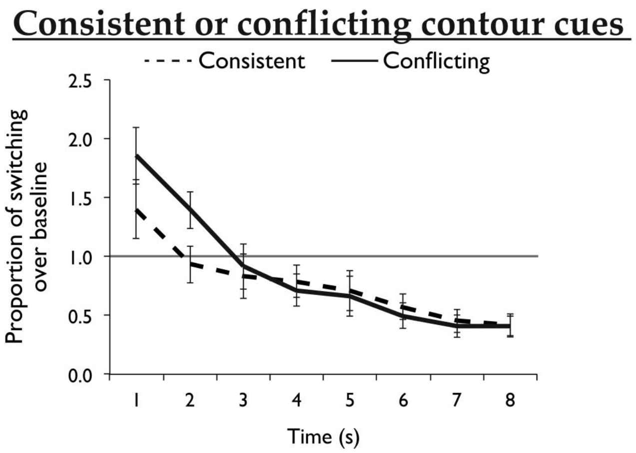

Figure 5. Observers' perceptual switching for Conflicting versus Consistent Cues in Experiment 4, showing that the initial above-baseline switching is enhanced for Conflicting Cues. All values are presented as proportions of switching above baseline (so that, e.g., a value of 2 indicates that there was twice as much switching relative to baseline). Error bars indicate 95% within-subject CIs.

to influence the rate of perceptual switching (Kanai et al., 2005). And even though most observers did not consciously perceive the transient cues in the present experiments, research has shown that attention can be readily captured and directed by cues (including transient cues) that are not seen (e.g., Lin et al., 2009; Norman et al., 2013). The effect with Consistent Cues may thus not reflect any theoretically novel processes. Accordingly, we focus our discussion below on the increased switching due in particular to the Conflicting Cues.

General Discussion

Perhaps the central fact of visual perception is just that there is too much incoming information to fully process to the level of conscious visual awareness. Accordingly, a bottleneck arises, and so there must exist processes that determine which stimuli do and do not reach this threshold. The features that make the difference may often be detected and processed unconsciously. This is presumably true for nearly all that we see, almost by definition: If a stimulus is not yet conscious, then the processes that help to promote it into awareness must themselves operate unconsciously. But this can also be appreciated in a more dramatic fashion, as in the case of perceptual "blindnesses." In motion-induced blindness, for example, salient attended targets may repeatedly vanish from awareness in the presence of a superimposed global motion pattern (Bonneh, Cooperman, & Sagi, 2001; New & Scholl, 2008). During the period in which a target has disappeared, however, a change can occur that will bring it back into awareness more quickly—for example, having it move (Mitroff & Scholl, 2004) or having it morph into a stimulus that is congruent with the contents of working memory (Chen & Scholl, 2013). Here we are asking related questions about visual awareness and unconscious processing, but with two key differences. First, we ask not only whether unconscious cues can promote a stimulus into awareness in a categorical sense, but also about whether they can change the specific content of what is seen—such that different kinds of cues will have different effects. Second, and most critically, we ask about whether unconscious cues can override our current conscious percept (during online perception) even when we are *never* aware of those cues, possibly throughout an entire experiment.

We tested these questions using the Spinning Dancer illusion. Like any bistable stimulus, the dancer can be seen in two ways—corresponding to two poses in three dimensions, and also two directions of rotation. We discovered that which interpretation is seen can be influenced by transient cues—the addition (and then immediate removal, after <100 ms) of subtle contours that disambiguate the dancer's orientation. This effect occurred even for extremely brief presentations, and critically it depended on the nature (and not just the existence) of the cue: Cues that conflicted with the observer's current percept led to increased perceptual switching shortly thereafter—more so than cues that were consistent with the observers' current percept. Most critically, this effect occurred even when observers failed to notice the cues themselves—not only on specific trials, but throughout the entire experiment.

This result adds a new twist to the broader literature on the power of unconscious stimuli to influence what we see. Of course, it is commonplace (in almost every instance of both illusory and veridical perception) for observers to remain unaware of just *how* various image cues are influencing their percepts. Take, for example, the perception of shape from shading, as in the impression of three-dimensional concavity or convexity due to the presence of lightness gradients (e.g., Ramachandran, 1988). Observers will readily experience such percepts, but they presumably have no idea just which image cues are causing them (viz., the lightness gradients), or why (viz., the operation of a light-from-above bias). Nevertheless, these observers are still able to readily see the lightness gradient itself, even if they fail to appreciate its importance. Researchers have also uncovered many instances in which what we see is influenced by visual cues that are themselves completely unseen (perhaps because they have been rendered invisible due to manipulations such as motion-induced blindness or continuous flash suppression). For example, unseen cues can change the perceived position of a subsequently presented object (e.g., Whitney, 2005), or can induce an aftereffect in a later independent stimulus (e.g., Kaunitz, Fracasso, & Melcher, 2011). Our results are another example of this sort of phenomenon, but with a new twist. Nearly all past studies of this type have involved an unseen cue influencing the perception of a later independent stimulus. In contrast, the current results demonstrate how an unseen transient cue can effectively *change* the nature of an already-existing, ongoing percept—in particular, by influencing the perception of an ambiguous figure during online perception. This is the first demonstration of this to our knowledge.

Though we screened participants for a lack of awareness of the transient contour cues in the main analyses, this lack of awareness was not especially uncommon: Throughout the three experiments with such cues presented here, 64.52% of observers failed to notice the cues at any time during the experiment. Critically, this value was very similar to that obtained in Experiment 1, where observers were given much more brute exposure to the cues, and were asked several types of questions to assess noticing. Of course, observers *could* have noticed the contours if they had been expecting them, but they apparently did *not* notice them when they were not expected. In essence, the observers were functionally blind to those cues: Even if they had seen them in some subtler, philosophical sense, they were unable to report them afterward on the basis of the types of questions typically used to assess inattentional blindness—the results of which have been shown to reflect deficits in perception rather than memory (Ward & Scholl, in press).

Thus unconscious cues can indeed change ongoing percepts even when we are unaware that those cues even exist.

A New Type of Cue to Switching?

Whereas previous work has routinely demonstrated that various stimulus factors can influence switching (e.g., Kanai et al., 2005; Klink et al., 2008), these manipulations were always overt—influencing what is seen, but also changing the content of what is seen. (Making some dots brighter in a structure-from-motion display might lead to a switch in perceived rotation direction, but it does so while also making the dots look brighter. Adding a sudden flash to a display might lead to a perceptual switch, but typically via the overt capture of attention—and awareness—by the flash.) In contrast, the present transient cues that led to switching were not noticed, as assessed by methods typical to the assessment of inattentional blindness. Similarly, whereas previous work has demonstrated that extrinsic transient cues can lead to perceptual switching (Kanai et al., 2005), the present results are the first to our knowledge that show how the specific content of those cues (here assessed by the presence of Consistent vs. Conflicting Cues) can have different effects. In our experiments, for example, one small quickly flashed line segment (that conflicts with your current percept) might cause you to switch what you see, while another very similar segment (consistent with your current percept) might have a much weaker effect.

One aspect of these results was unexpected, however: Though unnoticed transient cues led to perceptual switching shortly after the cues occurred, they also led to *below*-baseline switching for the later bins (see Figures 2 and 5)—and it remains uncertain just how this is to be explained. If this were only true for the Consistent Cues, then we might be tempted to explain it by appeal to some sort of strengthening of existing representations. But since this also occurred for the Conflicting Cues (in both Experiment 2 and Experiment 4), it may instead just reflect a more general factor—for example, of focused attention. Perhaps, for example, focused attention on certain features or locations makes a system less susceptible to switching that would otherwise arise from local stochastic network dynamics. If an interpretation of this latter sort were correct, then it would only strengthen the primary result of increased switching due to unconscious Conflicting Cues—since that effect might effectively be in conflict itself with such an attention-driven effect. Future work might be able to study such factors in a more focused way.

Conclusions: Implications for Visual Awareness

A central lesson of vision research (if not all of cognitive science) is that the bulk of the processing that determines our experience of the world is hidden from us, as it occurs unconsciously. Our phenomenology then masks this underlying complexity, so that we are aware of our seemingly "simple" percepts but not of the great underlying efforts that produced them. A similar dynamic can be seen in the present experiments, where what observers saw (in terms of the two possible interpretations of the Spinning Dancer) was determined at least in part by the specific character of the momentary cues. But critically, most observers (and all of them included in our analyses) failed to notice the existence of these cues in the first place. As a result, the ensuing perceptual switching (from the Conflicting Cues) was not only seemingly "simple," but it was seemingly random. In the present experiments, we were able to effectively change (in part) what our observers saw, but even when we specifically manipulated the content of the unconscious cues to generate a new percept, our observers experienced something apparently no different from people observing the Spinning Dancer animation its initial (fully ambiguous, silhouette) form—just a bunch of occasional haphazard switches. Thus, just as a phenomenological sense of ease should not be taken to imply a corresponding lack of underlying effortful computation, a phenomenological sense of randomness should not be taken to imply a corresponding lack of underlying systematicity.

References

Bonneh, Y. S., Cooperman, A., & Sagi, D. (2001). Motion-induced blindness in normal observers. *Nature, 411,* 798–801. http://dx.doi.org/10.1038/35081073

Borsellino, A., De Marco, A., Allazetta, A., Rinesi, S., & Bartolini, B. (1972). Reversal time distribution in the perception of visual ambiguous stimuli. *Kybernetik, 10,* 139–144. http://dx.doi.org/10.1007/BF00290512

Chen, H., & Scholl, B. J. (2013). Congruence with items held in visual working memory boosts invisible stimuli into awareness: Evidence from motion-induced blindness [Abstract]. *Journal of Vision, 13,* 808a. Retrieved from http://jov.arvojournals.org/article.aspx?articleid=2142915

Crump, M. J., McDonnell, J. V., & Gureckis, T. M. (2013). Evaluating Amazon's Mechanical Turk as a tool for experimental behavioral research. *PLoS ONE, 8,* e57410. http://dx.doi.org/10.1371/journal.pone.0057410

Elze, T. (2010). Misspecifications of stimulus presentation durations in experimental psychology: A systematic review of the psychophysics literature. *PLoS ONE, 5,* e12792. http://dx.doi.org/10.1371/journal.pone.0012792

Fox, R., & Herrmann, J. (1967). Stochastic properties of binocular rivalry alternations. *Perception & Psychophysics, 2,* 432–436. http://dx.doi.org/10.3758/BF03208783

Girgus, J. J., Rock, I., & Egatz, R. (1977). The effect of knowledge of reversibility on the reversibility of ambiguous figures. *Perception & Psychophysics, 22,* 550–556. http://dx.doi.org/10.3758/BF03198762

Gregory, R. L. (1980). Perceptions as hypotheses. *Philosophical Transactions of the Royal Society of London, Series B, 290,* 181–197.

Harrison, S. J., Backus, B. T., & Jain, A. (2011). Disambiguation of Necker cube rotation by monocular and binocular depth cues: Relative effectiveness for establishing long-term bias. *Vision Research, 51,* 978–986. http://dx.doi.org/10.1016/j.visres.2011.02.011

Helmholtz, H. (1866/1925). *Handbuch der physiologischen optik* [Treatise on physiological optics] (Vol. 3). (J. Southall, Trans.). Rochester, NY: Optical Society of America.

Jackson, S., & Blake, R. (2010). Neural integration of information specifying human structure from form, motion, and depth. *The Journal of Neuroscience, 30,* 838–848. http://dx.doi.org/10.1523/JNEUROSCI.3116-09.2010

Jiang, Y., Costello, P., Fang, F., Huang, M., & He, S. (2006). Gender and sexual orientation dependent attentional effect of invisible images. *Proceedings of the National Academy of Sciences of the United*

States of America, 103, 17048–17052. http://dx.doi.org/10.1073/pnas.0605678103

Kanai, R., Moradi, F., Shimojo, S., & Verstraten, F. A. (2005). Perceptual alternation induced by visual transients. *Perception, 34,* 803–822. http://dx.doi.org/10.1068/p5245

Kaunitz, L., Fracasso, A., & Melcher, D. (2011). Unseen complex motion is modulated by attention and generates a visible aftereffect. *Journal of Vision, 11,* 10. http://dx.doi.org/10.1167/11.13.10

Klink, P. C., van Ee, R., & van Wezel, R. J. (2008). General validity of Levelt's propositions reveals common computational mechanisms for visual rivalry. *PLoS ONE, 3,* e3473. http://dx.doi.org/10.1371/journal.pone.0003473

Kosegarten, J., & Kose, G. (2014). Seeing reversals in ambiguous images: To know or not to know? *Perceptual and Motor Skills, 119,* 228–236. http://dx.doi.org/10.2466/24.27.PMS.119c15z9

Leopold, D. A., & Logothetis, N. K. (1999). Multistable phenomena: Changing views in perception. *Trends in Cognitive Sciences, 3,* 254–264. http://dx.doi.org/10.1016/S1364-6613(99)01332-7

Levelt, W. J. M. (1966). The alternation process in binocular rivalry. *British Journal of Psychology, 57,* 225–238. http://dx.doi.org/10.1111/j.2044-8295.1966.tb01023.x

Lin, J. Y., Murray, S. O., & Boynton, G. M. (2009). Capture of attention to threatening stimuli without perceptual awareness. *Current Biology, 19,* 1118–1122. http://dx.doi.org/10.1016/j.cub.2009.05.021

Liu, C.-H., Tzeng, O. J. L., Hung, D. L., Tseng, P., & Juan, C.-H. (2012). Investigation of bistable perception with the "silhouette spinner": Sit still, spin the dancer with your will. *Vision Research, 60,* 34–39. http://dx.doi.org/10.1016/j.visres.2012.03.005

Long, G. M., & Toppino, T. C. (2004). Enduring interest in perceptual ambiguity: Alternating views of reversible figures. *Psychological Bulletin, 130,* 748–768. http://dx.doi.org/10.1037/0033-2909.130.5.748

Mack, A., & Rock, I. (1998). *Inattentional blindness*. Cambridge, MA: MIT Press.

Meredith, G. M., & Meredith, C. G. (1962). Effect of instructional conditions on rate of binocular rivalry. *Perceptual and Motor Skills, 15,* 655–664. http://dx.doi.org/10.2466/pms.1962.15.3.655

Mitroff, S. R., & Scholl, B. J. (2004). Seeing the disappearance of unseen objects. *Perception, 33,* 1267–1273. http://dx.doi.org/10.1068/p5341no

Mitroff, S. R., Sobel, D. M., & Gopnik, A. (2006). Reversing how to think about ambiguous figure reversals: Spontaneous alternating by uninformed observers. *Perception, 35,* 709–715. http://dx.doi.org/10.1068/p5520

Moreno-Bote, R., Shpiro, A., Rinzel, J., & Rubin, N. (2010). Alternation rate in perceptual bistability is maximal at and symmetric around equidominance. *Journal of Vision, 10,* 1–10. http://dx.doi.org/10.1167/10.11.1

Most, S. B., Scholl, B. J., Clifford, E. R., & Simons, D. J. (2005). What you see is what you set: Sustained inattentional blindness and the capture of awareness. *Psychological Review, 112,* 217–242. http://dx.doi.org/10.1037/0033-295X.112.1.217

New, J. J., & Scholl, B. J. (2008). "Perceptual scotomas": A functional account of motion-induced blindness. *Psychological Science, 19,* 653–659. http://dx.doi.org/10.1111/j.1467-9280.2008.02139.x

Norman, L. J., Heywood, C. A., & Kentridge, R. W. (2013). Object-based attention without awareness. *Psychological Science, 24,* 836–843. http://dx.doi.org/10.1177/0956797612461449

Pastukhov, A., Vonau, V., & Braun, J. (2012). Believable change: Bistable reversals are governed by physical plausibility. *Journal of Vision, 12,* article 1.

Peirce, J. W. (2007). PsychoPy—Psychophysics software in Python. *Journal of Neuroscience Methods, 162,* 8–13. http://dx.doi.org/10.1016/j.jneumeth.2006.11.017

Peterson, M. A., & Gibson, B. S. (1991a). Directing spatial attention within an object: Altering the functional equivalence of shape descriptions. *Journal of Experimental Psychology: Human Perception and Performance, 17,* 170–182. http://dx.doi.org/10.1037/0096-1523.17.1.170

Peterson, M. A., & Gibson, B. S. (1991b). The initial identification of figure-ground relationships: Contributions from shape recognition processes. *Bulletin of the Psychonomic Society, 29,* 199–202. http://dx.doi.org/10.3758/BF03342677

Peterson, M. A., & Gibson, B. S. (1994). Object recognition contributions to figure-ground organization: Operations on outlines and subjective contours. *Perception & Psychophysics, 56,* 551–564. http://dx.doi.org/10.3758/BF03206951

Ramachandran, V. S. (1988). Perception of shape from shading. *Nature, 331,* 163–166. http://dx.doi.org/10.1038/331163a0

Rock, I. (1983). *The logic of perception*. MIT Press, Cambridge, MA.

Rock, I., & Mitchener, K. (1992). Further evidence of failure of reversal of ambiguous figures by uninformed subjects. *Perception, 21,* 39–45. http://dx.doi.org/10.1068/p210039

Taylor, M. M., & Aldridge, K. D. (1974). Stochastic processes in reversing figure perception. *Perception & Psychophysics, 16,* 9–25. http://dx.doi.org/10.3758/BF03203243

Tong, F., Meng, M., & Blake, R. (2006). Neural bases of binocular rivalry. *Trends in Cognitive Sciences, 10,* 502–511. http://dx.doi.org/10.1016/j.tics.2006.09.003

Toppino, T. C. (2003). Reversible-figure perception: Mechanisms of intentional control. *Perception & Psychophysics, 65,* 1285–1295. http://dx.doi.org/10.3758/BF03194852

Troje, N. F. (in press). The Kayahara silhouette illusion. In A. Shapiro & D. Todorovic (Eds.), *Oxford compendium of visual illusions*.

Troje, N. F., & McAdam, M. (2010). The viewing-from-above bias and the silhouette illusion. *i-Perception, 1,* 143–148. http://dx.doi.org/10.1068/i0408

Walker, P. (1975). Stochastic properties of binocular rivalry alternations. *Perception & Psychophysics, 18,* 467–473. http://dx.doi.org/10.3758/BF03204122

Ward, E. J., & Scholl, B. J. (in press). Inattentional blindness reflects limitations on perception, not memory: Evidence from repeated failures of awareness. *Psychonomic Bulletin & Review*.

Whitney, D. (2005). Motion distorts perceived position without awareness of motion. *Current Biology, 15,* R324–R326. http://dx.doi.org/10.1016/j.cub.2005.04.043

Windmann, S., Wehrmann, M., Calabrese, P., & Güntürkün, O. (2006). Role of the prefrontal cortex in attentional control over bistable vision. *Journal of Cognitive Neuroscience, 18,* 456–471. http://dx.doi.org/10.1162/jocn.2006.18.3.456

Yu, K., & Blake, R. (1992). Do recognizable figures enjoy an advantage in binocular rivalry? *Journal of Experimental Psychology: Human Perception and Performance, 18,* 1158–1173. http://dx.doi.org/10.1037/0096-1523.18.4.1158

Received May 6, 2014
Revision received October 8, 2014
Accepted October 9, 2014 ■