WE DON'T NEED TO KNOW WHAT WE SEE: MODEST MEDIATION OF BISTABLE PERCEPTION BY KNOWLEDGE

By

Bobicheng Zhang

A THESIS

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Psychology – Master of Arts

2022

ABSTRACT

WE DON'T NEED TO KNOW WHAT WE SEE:
MODEST MEDIATION OF BISTABLE PERCEPTION BY KNOWLEDGE

By

Bobicheng Zhang

When faced with ambiguous visual input, an observer may have various perceptual interpretations even when the input remains unchanged. Indeed, such ambiguous displays can cause the observer to experience distinct perceptual interpretations in turn, unpredictably switching between them over time. Theories of such so-called multistable perception broadly fall into two categories: top-down theories that hold that perception in these cases depends strongly on higher-level cognitive factors such as knowledge, and bottom-up theories which suggests more vital involvement of aspects of lower-order information processing such as local adaptation in the visual system. We evaluated whether the occurrence of perceptual reversals in the face of ambiguous input is related to the observer's knowledge that the input is, indeed, ambiguous. We used an ambiguous animation that was designed such that subjects could report perceptual reversals without realizing the ambiguity. Subjects observed the animation, reported their perception, and filled out a questionnaire that assessed their knowledge of the animation's ambiguity. We found that informed subjects reported slightly more perceptual switches than the other subjects, but that this between-group difference was very small compared to the lack of variability within each group between subjects who were aware of the ambiguity and those who were not. These findings suggests that knowledge of ambiguity can influence perception of ambiguous stimuli, but that this influence is relatively minor. This discrepancy between current findings and past work is discussed.

ACKNOWLEDGEMENTS

Throughout this project I have received a great deal of support.

I would first like to thank my advisor, Dr. Jan Brascamp, whose expertise was invaluable in formulating the research questions and methodology. I would also like to thank the rest of my Master's committee, Dr. Mark Becker and Dr. Taosheng Liu for their feedback and guidance during the early stages of the project.

I would like to acknowledge the other members of the Brascamp Lab, Haley Frey and Abhilasha Jagtap for their assistance and interesting discussions throughout the writing of this thesis.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

In everyday visual perception, our brain is spoiled with relatively unambiguous information that typically has one overwhelmingly plausible interpretation. However, when exposed to ambiguous information that is equally consistent with more than just a single percept, the visual system responds by favoring each of these perceptual interpretations in turn. That is, such ambiguous displays typically lead the viewer to perceive one coherent interpretation at a time yet to switch between interpretations unpredictably. This is because both interpretations of the display are consistent with the input, and both are likely to be extrapolated based on the available information, yet in reality they are deemed impossible to co-exist. For example, a quintessential ambiguous figure, the Necker Cube can be seen as facing in two different directions (Fig. 1.1). It is ambiguous because it creates an impossible scenario for the visual system with two cube-like shapes that appear to occupy the same physical space. The conflict between perceiving the downward facing cube and the upward facing cube drives perceptual reversals, as the observer tends to reverse between the two randomly.

A variety of theories has been proposed to understand the underlying processes leading to perceptual reversals, yet contributing factors are not clearly understood. Many of those factors align with two historically popular classes of theories for bistable perception: bottom-up theories and top-down theories. The bottom-up theories such as Hering's sensory adaptation hypothesis Hering (1964), consider perceptual transitions to be the result of passive, automatic sensory process that are relatively independent of cognitive control. Applying this principle to the case of the Necker Cube suggests that reversals between the two cubes are stimulus-driven and the observer cannot voluntarily alter their interpretation. On the other hand, the top-down perspectives such as Helmholtz's attentional shift hypothesis Helmholtz (1962), consider them to be active, cognitive processes wherein the observer retains more control that can influence their perception. In the case of the Necker Cube, this would mean that observer has some degree of control over which orientation is perceived, downward or upward facing.

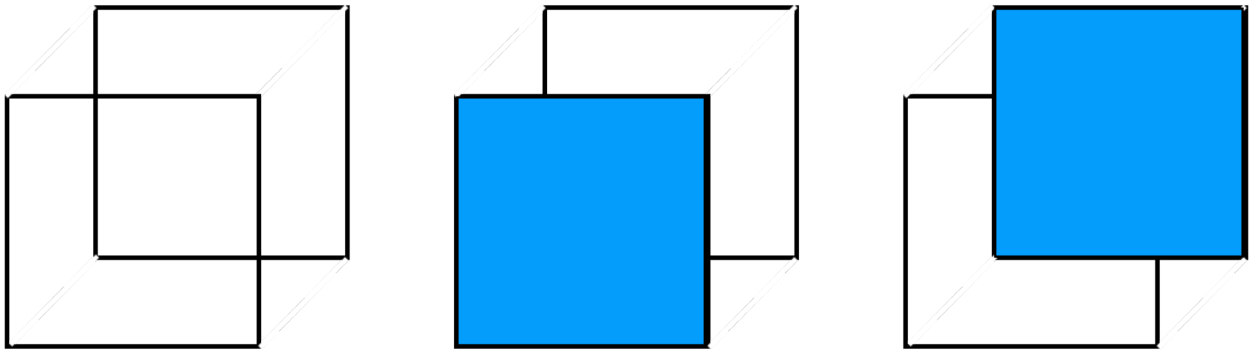

Figure 1.1: The Necker Cube and its two perceptual alternatives. It can be seen as downward-facing or upward-facing, and typically the observer switches between seeing one and the other.

The same broad distinction can still be identified in more modern theories on perceptual alternations (Girgus et al., 1977). In recent times, it has been customary to view perceptual reversals as a combined product of both top-down and bottom-up factors. In the current study however, we focus on one specific top-down factor that may influence the perception of ambiguous displays, namely explicit knowledge about the ambiguous nature of what is being viewed. It has been theorized that the ability to reverse ambiguous displays, or switch between multiple perceptual interpretations, necessitates explicit knowledge because awareness of the reversibility of an ambiguous object enables the observer to create an association between the percepts of the object and the object itself (Ach, 1910; Girgus et al., 1977; Lewin, 1922). Such associations were speculated to be essential: understanding the ambiguity of the image would allow the observer to experience perceptual reversals. Consistent with this idea, empirical studies using a range of ambiguous figures have reported that knowledge of stimulus ambiguity played an important role in ensuring that both versions of the figures were perceived. Observers who were informed of the reversibility of the figures experienced perceptual reversals more quickly and more frequently than those who were oblivious (Girgus et al., 1977; Rock & Mitchener, 1992). For instance, one study showed that in a within-subjects design, at least half of all subjects reported seeing no reversals at all before learning about the reversibility of the figures they had just viewed. Yet after they were informed that the figures had been ambiguous and could be perceived in two ways, all subjects were able to switch between perceptual interpretations upon their next viewing (Girgus et al., 1977).

Numerous subsequent studies all consistently showed a beneficial impact of knowledge in aiding observers switch between multiple perceptual interpretations of ambiguous objects (Girgus et al., 1977; Kosegarten & Kose, 2014; Rock et al., 1994; Rock & Mitchener, 1992). Not only was knowledge shown to facilitate the occurrence of perceptual reversals for those who would otherwise not reverse, it drastically increased the number of reversals. For instance, in another within-subject study using drawings of ambiguous figures, most subjects advanced from reporting no reversal at all to reversing once every 4 to 5 seconds during a 30-second viewing period once they became informed of the possible perceptual alternatives (Rock & Mitchener, 1992). This trend was replicated in a later study with a between-subjects design, although the authors of that study did point out that knowledge is no guarantee for reversal (Kosegarten & Kose, 2014).

Additionally, knowledge was also reported to contribute to the occurrence of perceptual reversals in the sense that it reduced the time taken to achieve the initial reversal after stimulus appearance, such that informed observers showed a tendency to experience their first reversal earlier than the uninformed observers. In a series of experiments during which ambiguous figures were presented for 30, 60, 120 and 180 seconds, at least 80% of the informed subjects, in contrast to approximately half of the uninformed subjects, consistently reversed within the first 30 seconds of inspection across all experiments (Girgus et al., 1977). Although not directly reported by Rock and Mitchener (1992), the near total lack of reversals among their uninformed observers, and the multiple reversals reported by most informed observers, are also consistent with an earlier time for initial reversal for the informed observers. After all, almost all uninformed participants, apparently, needed more than the 30-second viewing time to experience their first reversal.

While these findings are consistent with the notion that knowledge facilitates the occurrence of perceptual reversals, no definitive account as to why knowledge would have this effect has been offered. These authors have generally argued that ambiguous displays are not reversible to the observer and do not trigger perceptual reversals unless the observer knows that they are indeed reversible (Kosegarten & Kose, 2014; Rock & Mitchener, 1992). They have speculated regarding the underlying mechanism that the explicit knowledge of a figure's ambiguity seems to lead to an involuntary process facilitating perceptual reversals. This was sometimes referred to as a type of intention that leads to spontaneous reversals. Such intention could be built upon the foundation that the observer already has memory traces of the perceptual alternatives of a figure, and further established an association between them (Leopold & Logothetis, 1999; Lewin, 1922; Rock et al., 1994). Overall, no conclusive account of the mechanisms behind the apparent effect of knowledge has been identified.

Our reason for revisiting the role of knowledge was that many currently popular accounts of bistable perception have a substantial bottom-up flavor. In particular, a predominant view of perceptual reversals is that these reversals are in considerable part due to gradual adaptation of sensory representations associated with the currently-perceived interpretation (Blake et al., 2003; Chen & He, 2004; Leopold & Logothetis, 1999; Pastukhov et al., 2014). It is rare these days to find any account that could be classified purely bottom-up or purely top-down, and even adherents of an adaptation-centered viewpoint are likely to agree that factors like, say, attention to particular parts of the stimulus or an intention to hold on to a particular perception can influence the timing and frequency of perceptual reversals (Chong et al., 2005; Meng & Tong, 2004; van Ee et al., 2005). Still, the above-cited literature on the impact of knowledge does not sit well with adaptation-centered viewpoints: if perceptual reversals arise due to sensory adaptation, typically thought to be automatic, of the dominant representation, then a modest modulatory influence of top-down factors is conceivable (and is generally what studies of attention and intention report), but the literature on knowledge shows findings that go well beyond that. As outlined above, these findings include extreme ones like a wholesale lack of reversals in a majority of uninformed participants and increases in reversal rate from fewer than one in half a minute to at least six times that rate. In sum, then, one reason for us to re-examine the role of knowledge is this apparently unresolved conflict between, on the one hand, earlier reports of a major dependence of reversals on knowledge and, on the other, the popular view throughout recent literature as a whole that perceptual switching should to a large degree unfold automatically.

Another important reason for revisiting the role of knowledge, is that there are some limitations to the previous work that examined the topic. For one, as was pointed out by Kosegarten and Kose Kosegarten and Kose (2014), most previous work may have overlooked the possibility of priming. By 'priming' in this context we mean the possibility that perceptual reversals occur more frequently in a visual system in which they have arisen before (Kosegarten & Kose, 2014; Rock & Mitchener, 1992). It is natural that most studies on the role of knowledge, presumably aiming for knowledge level to be the only independent variable, have used a within-subjects design in which subjects first reported their perceptual experiences in an uninformed condition, were then told about the ambiguous nature of the figures, and then again reported their perceptual experiences, now in an informed condition (Girgus et al., 1977; Rock & Mitchener, 1992). Nevertheless, this means that, as participants moved on to the informed condition from the uniformed condition, in addition to learning about the ambiguity of the figures from the experimenters, they could have also been primed by having seen the same ambiguous figures prior to being given the confirmation. Consistent with this idea, most reversals reported in the uninformed condition in the work by Rock & Mitchener Rock and Mitchener (1992) occurred at the end of the viewing period, and their frequency increased with exposure time, indicating a possible role of priming. The use of such sequential within-subject designs in existing work may be particularly relevant given the idea that an observer's first perceptual reversal may have a particularly strong reliance on knowledge, as compared to subsequent reversals (Rock et al., 1994).

A second potential limitation of existing work is that it may not have been successful at creating situations in which subjects were fully naïve regarding the ambiguity of the display (Girgus et al., 1977; Rock & Mitchener, 1992). By 'naïve' here we mean that observers without explicit prior knowledge of perceptual ambiguity would have no suspicion whatsoever that the figure in front of them was worth further examination and re-evaluation. The tasks used in those earlier studies were such that their instructions or procedures might have invited such a suspicion anyway. For instance, in some cases observers were asked to repeatedly report what they perceived in a line drawing on paper in front of them whenever they detected a change (Girgus et al., 1977), or at fixed intervals every several seconds (Rock & Mitchener, 1992). Such instructions imply that the observer's perception of the drawing may change following their initial report, even if the observer has not been explicitly told that the drawing shows an ambiguous figure, or what the possible perceptions are. Furthermore, as Kosegarten and Kose noted in their report Kosegarten and Kose (2014), the experimental instructions often included potentially suggestive phrases, such as "what you can see in the image," that could be interpreted by the subjects to mean that there were alternative percepts. If such experiment features invited additional scrutiny on the part of the participant in the supposedly uninformed conditions, then the number of perceptual reversals recorded in such conditions, although reduced relative to informed conditions, might still present an overestimation. In that case, the real impact of knowledge may be even larger than estimated in those studies.

The current experiment is designed to re-evaluate how the occurrence of perceptual reversals is related to the observer's knowledge of stimulus ambiguity, but now using a between-participant design to avoid issues of priming and using a specific stimulus and task that should minimize suspicion of ambiguity in uninformed participants. In particular, we manipulate knowledge of stimulus ambiguity by informing half of the participants about the ambiguity of our stimulus prior to the experiment, and giving the other participants a different instruction prior to the same experiment. In both conditions, participants report perception of an ambiguous animation of a sphere that undergoes two types of gradual motion: one that is vital to its perceptual reversibility, and one that has nothing to do with the ambiguous nature of the stimulus. Critically, the task requires that participants report an aspect of their perception that changes whenever perception reverses due to the ambiguity, but that also changes because of the on-screen motion that is unrelated to ambiguity. In other words, a participant's response will change periodically irrespective of whether the participant experiences perceptual reversals. In this way we circumvent the issue described above, namely that in some existing work the task itself could only make sense to participants if they assumed there was more to the stimulus than met the eye.

Aside from measuring how participants in the informed condition and in the uninformed condition differ in their reported perceptual experiences, we also aimed to measure how successful we had been at informing our participants of the ambiguity and, more generally, how much relevant knowledge our participants had to begin with. Thus, like previous studies that conducted individual interviews following the inspection of ambiguous figures (Girgus et al., 1977; Rock & Mitchener, 1992), we included a post-experiment survey that assessed these quantities. Note that, while the interviews in existing work were designed to also capture reversals that were experienced by the participants but not reported, our survey did not have this purpose. Instead, it only contained short forced-choice questions aimed at determining knowledge and prior experience regarding ambiguous displays.

If findings from existing literature hold in these arguably more controlled conditions and for this stimulus, then we expect to find a substantial number of perceptual reversals among those who were aware of the sphere's ambiguous nature at the time they viewed the sphere, and a drastically reduced number among those who were not. In addition, in that scenario we also expect knowledgeable participants to take less time before reporting their first perceptual reversal than their uninformed counterparts.

CHAPTER 2

MATERIALS AND METHODS

2.1 Participants

A total of 800 subjects were recruited for this study. All of them were undergraduate students at Michigan State University who registered for the study on the HPR/SONA system in compensation for psychology course credit. Informed consent was obtained from all participants prior to the experiment and a debriefing form was provided at the end of the study.

2.2 Stimulus

The current study used an ambiguous rotating sphere that could be perceived as rotating in one direction at times and the opposite direction at other times (Fig. 2.1A & B). The sphere has an on-screen diameter of 4 cm and consists of 600 randomly placed slow-moving white dots. This was an online study (see below) and we did not attempt to control viewing distance. However, assuming viewing distances within the range that is typical in such settings (Brascamp, 2021) we estimate the angular diameter of the stimulus to lie between 3.8 and 4.5 degrees. The rotation speed was maintained at 0.2 cycles per second. When the stimulus is positioned upright, half of all dots move toward the left while the other half move toward the right. The dots make up a transparent sphere of which both the front surface moving in one direction and the back surface moving in the opposite direction are simultaneously visible. However, it is ambiguous to observers which surface of the sphere is the front surface and which is the back surface, making the perceived direction of rotation fluctuate over time in many existing studies.

In addition to the sphere's perceived rotation direction potentially changing over time (blue arrows in Figs 2.1A and 2.1B), which is standard for this type of stimulus, we added a further change-over-time to this stimulus: in our paradigm the sphere also performs orientation changes in the plane of the screen (Fig. 2.1C). Unlike the rotation depicted in Figs 2.1A and 2.1B, these latter orientation changes appear entirely two-dimensional. And in contrast to the reversals in rotation direction that some observers perceive, which are endogenously generated and abrupt, these orientation changes occur on the screen and gradually, either in a clockwise or a counterclockwise fashion. Throughout the task, all participants are presented with the same sequence of gradual orientation changes, but this sequence does differ from trial to trial.

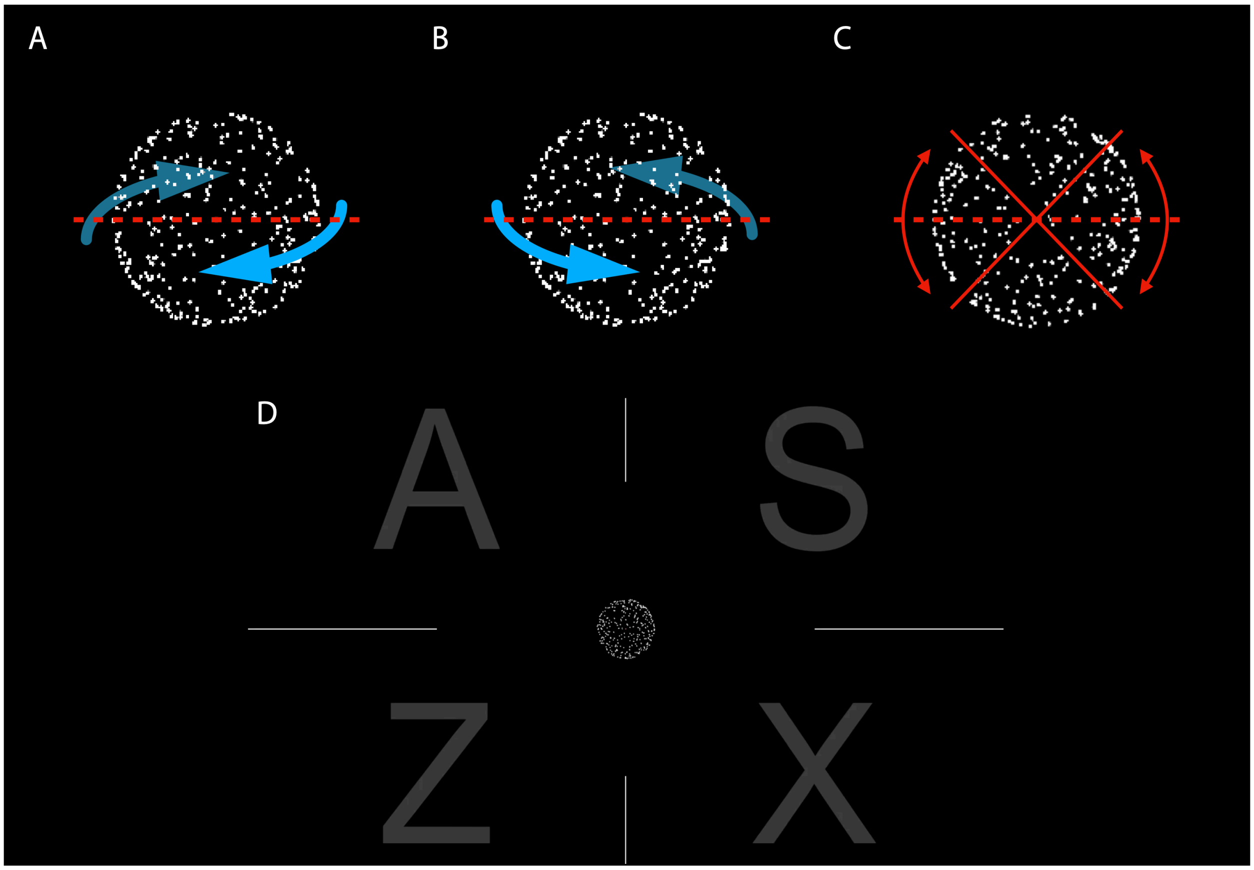

Figure 2.1: Schematics of the rotating sphere. A and B demonstrate examples of how the sphere may look at the same moment. The blue arrows indicate the perceived front and the back surface and are not present during the experiment. The red dashed lines indicate the sphere's equator, and are added here to depict the sphere's orientation. As the sphere is positioned in panels A and B it could be perceived as rotating to the left or rotating to the right as either surface could be seen as the front surface. C depicts gradual orientation changes and the range of orientation the sphere may move to. As shown here, the orientation of the sphere's equator stays within -45 to 45 degrees. D shows the screen as a participant would see during the task. The screen is divided into four quadrants with each labeled by a letter corresponding to the key they should press when they see the sphere rotate toward that quadrant.

2.3 Procedure

The experiment was programmed using PsychoPy (version 2021.1.4) and hosted on Pavlovia. Subjects performed the task on their own computer. To compensate for variations in monitor resolution and scaling, a procedure was adopted that allowed each individual participant to adjust a rectangle on screen to match the size of a credit card (Brascamp, 2021; Li et al., 2020) (adapted from gitlab.pavlovia.org/Wake/screenscale). This procedure modified scaling of the sphere stimulus to ensure both the correct aspect ratio and equal on-screen size across participants. During the experiment, participants were instructed to pay attention to the rotating sphere located at the center of the screen. The orientation of the sphere (or, equivalently, the orientation of its equator) started at a random angle (between -45 and 45 deg on the right, Fig. 2.1C) and unpredictably changed while the sphere remained at the center of the screen and the dots continued to move. Also on the screen are four lines that divided the screen into four quadrants like those of the Cartesian coordinate system (an xy-plane, Fig. 2.1D), with the sphere in its center. Participants were asked to indicate using keyboard buttons which quadrant the dots on the front of the sphere are moving toward, and to press a different key whenever their answer changes. Critically, the response to this question depended on the sphere's perceived rotation direction but also on its orientation, so while the task provided the experimenters with information on any changes in perceived rotation direction, the task was also meaningful to any participants who do not experience such changes.

In the uninformed condition, participants were first shown a static schematic of the sphere stimulus (similar to Fig. 2.1A and 2.1B) and a series of descriptions of the animated sphere so that they could understand the sequence of movements performed by the sphere without seeing the actual animation prior to the practice trial. This helped minimize their exposure to the stimulus. They were told that the sphere during the experiment would rotate as suggested by the schematic. Notably, it was never stated in this condition that the sphere was ambiguous or that one and the same sphere could be perceived as rotating in two directions. Nor was it stated that rotation direction could reverse (irrespective of whether reversals are endogenous). However, it was clearly explained that the sphere would continuously change its orientation. Participants were then instructed to pay close attention to the sphere at the center of the screen and indicate the quadrant to which the sphere's front surface dots were moving.

In contrast, participants in the informed condition were shown the sphere animation from the beginning of the experiment, including during their instructions. Importantly, they were told that the animation was consistent with a sphere rotating to one direction but also with a sphere rotating to the opposite direction. This means that participants in this condition, if they understood the instructions, were aware that the sphere animation was ambiguous during the experiment. Further, they were provided with more details pertaining to the two distinct types of change in direction that they might see. Specifically, it was made clear to them that the perceived direction of rotation could fluctuate because the sphere was ambiguous, and that the orientation could fluctuate on the screen. They were then given the same instructions as participants in the uninformed condition regarding the task and response keys, and proceeded to the rest of the experiment in the same manner as participants in the uninformed condition.

Participants in both conditions viewed the rotating sphere stimulus and reported their perceived direction of rotation for a total of 180 seconds (30*6 trials). After the main task was completed, each participant was also asked to complete a short survey in which they indicate their prior experience with ambiguous displays and their degree of knowledge during the sphere task. The survey contained 6 two-choice questions (Appendix A).

2.4 Data analysis

Due to the relatively uncontrolled nature of remotely hosted studies, a few steps were taken prior to our main data analysis to ensure the validity and reliability of results. First, participants whose data suggested that they failed to properly follow the procedure to adjust the size and aspect ratio of the stimulus were excluded. As previously mentioned, participants performed a stimulus scaling procedure to ensure that the subsequent sphere would be displayed at the intended size and aspect ratio. This involved the participant independently scaling the horizontal and vertical dimension of an on-screen stimulus, and this scaling procedure produced measures of the horizontal as well as the vertical pixel size of the participant's screen. We excluded participants whose scaling procedure indicated a pixel aspect ratio that deviated more than 5% from perfectly square. Because previous work indicated that almost all online participants have screens with square pixels (Brascamp, 2021), such deviations suggest insufficient dedication to the scaling procedure. The second step was to eliminate those participants who responded 'no' to a catch question in the post-experiment questionnaire (Q4, supplementary material). This question asked whether the participant had noticed the gradual in-plane orientation changes of the sphere, which were objectively present on the screen throughout the experiment, and which were also announced to participants during the instruction phase in both conditions. Third, participants who took too much time before their first key press of the individual trials were also excluded. This key press was supposed to signal their initial assessment of the quadrant toward which the sphere was rotating at the start of a trial. As the sphere was designed to always change its orientation within the first 5 seconds of each trial, those who took more than 5 seconds to report the initial quadrant, averaged across all trials, were excluded from analysis. Lastly, participants with subpar task performance were also excluded regardless of their perceptual reversals. This was possible because the reporting task involving the four quadrants had right and wrong responses throughout each trial: at any moment, irrespective of perceived rotation direction, there were only two quadrants toward which the participant could possibly be seeing the sphere rotating (see Fig. 2.2). Accordingly, we calculated for each participant the proportion of time the indicated quadrant was among the two possible ones and excluded participants for whom this proportion was lower than 0.6. Please see Fig. 2.2 for an example of a trial that corresponds to a high score on this measure (top panel) as well as an example corresponding to a low score (bottom panel; see figure label and next paragraph for further information on the figure's formatting conventions).

Our main measures from the sphere task were the number of perceptual reversals reported by each participant per unit time, and the amount of time it took each participant to report their first perceptual reversal. To accurately extract the number of perceptual reversals for each participant, we did not simply count the number of reported changes in the quadrant toward the sphere's front

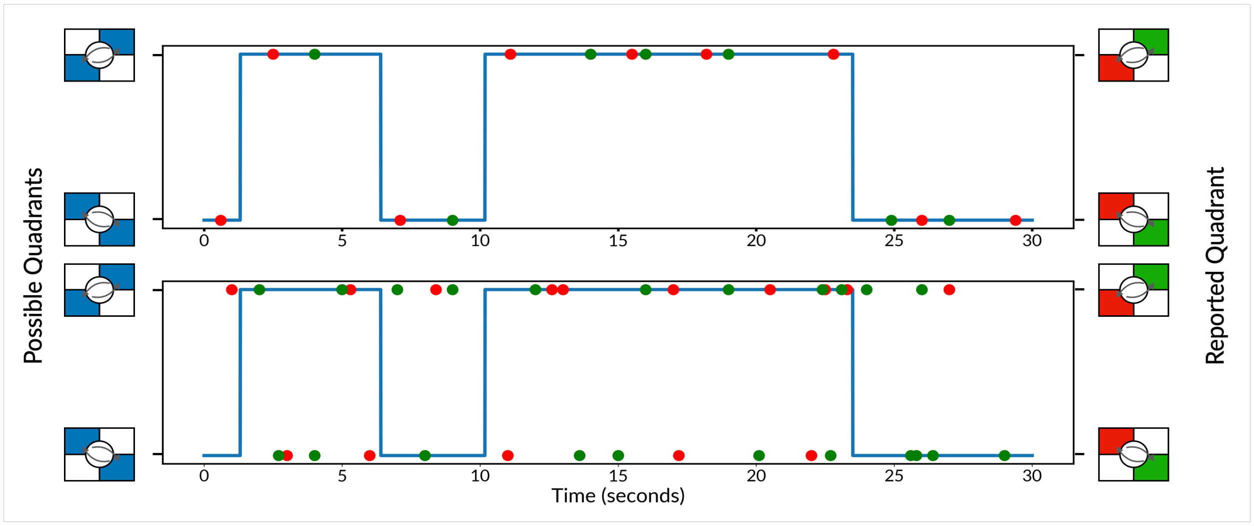

Figure 2.2: Examples of a good trial (top) and a bad trial (bottom). Depending on the on-screen orientation of the sphere (see the red dashed lines in Fig. 2.1), there are only two quadrants toward which it may be perceived to rotate at any given time. These 'possible' quadrants are indicated along the left y-axis and mapped out over time using blue lines. The quadrants that the participant reported are represented by red and green dots. Dot color and vertical position of a dot combined designate which of the four quadrants was reported (right y-axis). Because of this plotting choice, dots that fall on the blue line indicate valid responses (the participant reports a quadrant toward which the sphere may, indeed, be seen to rotate at that time), and dots that fall off the blue line indicate invalid responses. The horizontal axes represent time in the trial.

surface was rotating. Instead, we distinguished between changes that were due to the on-screen orientation changes of the sphere, and changes that indicated a reversal in perceived rotation direction. In terms of the schematic introduced in Fig. 2.2, a change in color from one dot to the next dot along the x-axis, irrespective of its position on the y-axis, was counted as a perceptual reversal of rotation direction, whereas a change in y-axis position without a color change was simply a reaction to the in-plane orientation change. For example, for the sample trials in Fig. 2.2 our algorithm would count 14 (top) and 21 (bottom) reversals, respectively (Although note that the data from the latter trial would likely be discarded due to that small proportion of time occupied by valid reported quadrants; see above). For our measure of the amount of time until the first perceptual reversal, we counted time cumulatively from the start of the real trial until the first reversal. In other words, if the first reversal did not occur until a later trial, then we added up time across trials.

CHAPTER 3

RESULTS

For comprehensive descriptive statistics, see Appendix C. After all exclusion criteria were applied, a total of 472 participants were included in the analyses, with 229 participants in the uninformed group and 243 participants in the informed group. All participants reported at least one perceptual reversal. On average over the entire duration of the sphere task, the uninformed participants reported 32.9 reversals (*SD* = 22.20) at a rate of 0.18 reversals per second (*SD* = 0.12) and the informed participants reported 39.9 reversals (*SD* = 36.10) at a rate of 0.22 reversals per second (*SD* = 0.20). Mann-Whitney U tests were performed to test the null hypothesis that the rate of perceptual reversals does not differ between groups (McKnight & Najab, 2010). This null hypothesis was rejected: between the informed and uninformed condition, significantly more perceptual reversals were reported among the informed participants (Fig. 3.1A, $p = .003$). This is consistent with the idea that participants more readily perceive abrupt direction reversals when they were informed of the ambiguity of the stimulus. In existing work, the impact of knowledge on reversal rate was very large which, as outlined in the introduction, motivated this study. In our study, the impact appears to be smaller. In particular, despite a significant role of knowledge, we found considerable overlap in across-participant distributions of perceptual reversal rate (Fig. 3.2A). We quantified the size of the effect of knowledge on reversal rate by creating a ROC curve (receiver operation characteristic, Fig, 3.2B). The area under the ROC curve is approximately 0.579. Given that an area of 0.5 corresponds to no effect of knowledge, whereas an area of 1 corresponds to an effect of knowledge that overshadows all between-subject differences in reversal rate, this analysis shows that the effect of knowledge on reversal rate is very modest in our data.

With regard to the time taken before reporting the first perceptual reversal, on average, the uninformed participants took 10.0 seconds (SD = 12.3) and the informed participants took 8.76 seconds (SD = 8.56). Although the direction of this trend was consistent with the direction of the effect on reversal rate, there was no significant difference in time until first reversal between the

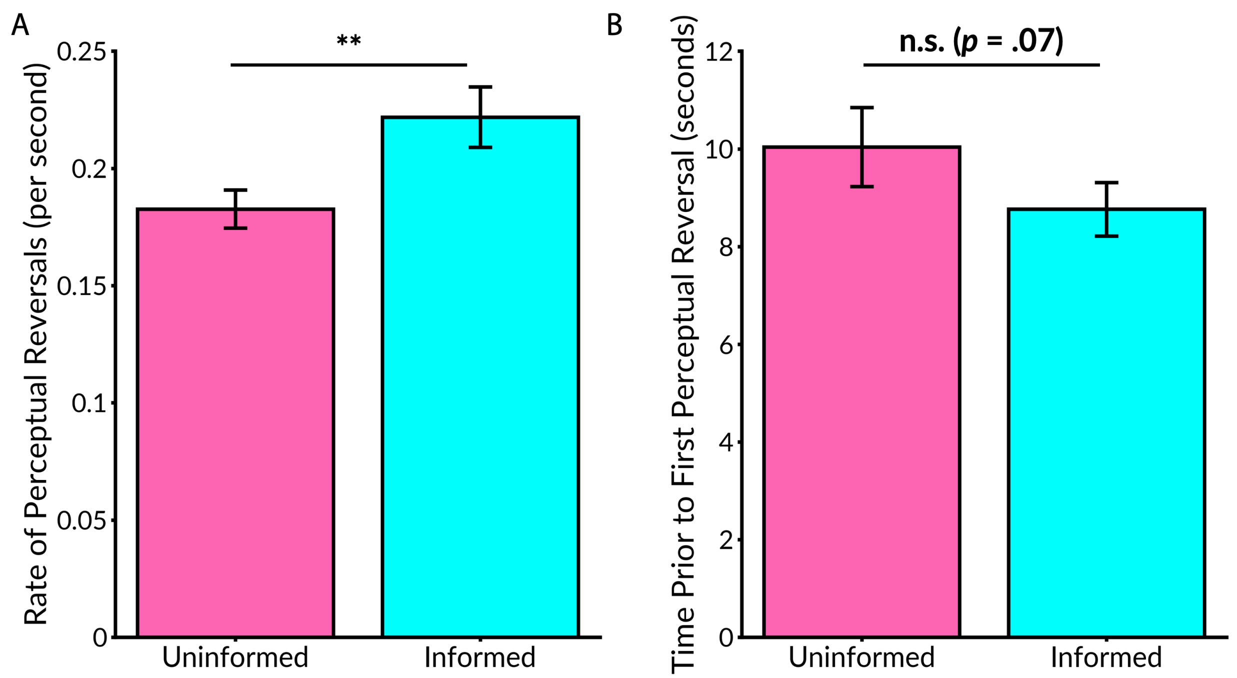

Figure 3.1: Between-conditions results. The left bars represent the average rate of perceptual reversals per second. Subjects in the informed condition reported significantly more perceptual reversals than those in the uninformed condition. The right bars represent the average time in seconds taken prior to participants' initial reversal. No difference was found in this comparison. * $p < .05$, ** $p < .01$, *** $p < .001$

uninformed and the informed group (Fig. 3.1B, p = .07).

Aside from our main analyses regarding the difference in perceptual reversal dynamics between conditions, which were designed to differ in terms of participant knowledge, we also considered participant knowledge as measured using our survey. First, for individual survey questions a Fisher's exact test was performed comparing responses between the two conditions. Here, we focus on question 6, which indexed whether observers thought that the sphere was ambiguous during the experiment (Fig. 3.3). Results for other survey questions are reviewed later.

Compared to the uninformed condition, a larger proportion of subjects in the informed condition reported thinking, at the time they were performing the rotation direction task, that the sphere was ambiguous (Fig. 3.3, $p = .003$). This result implies that our between-condition manipulation was effective. Nevertheless, numerous participants in the informed condition reported not thinking

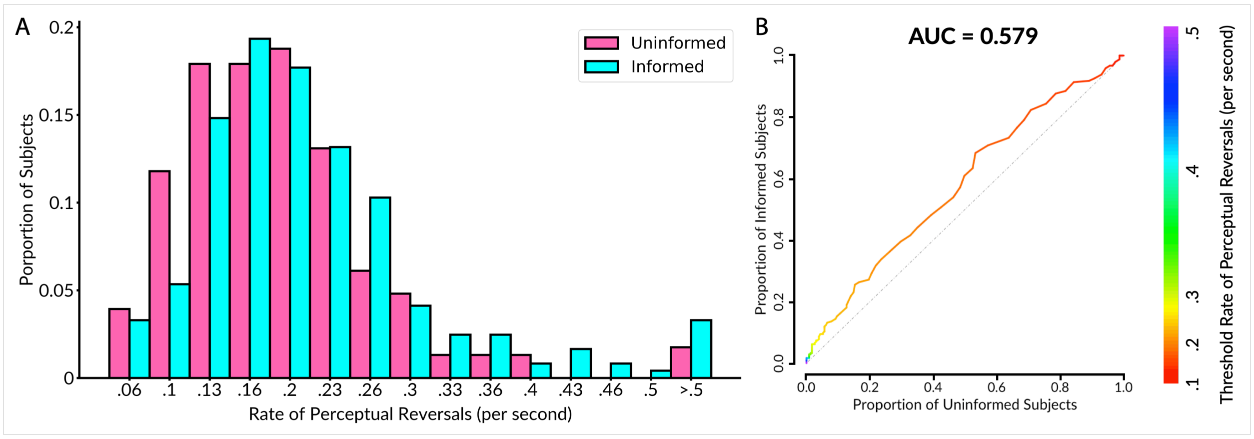

Figure 3.2: Between-conditions distributions of the rate of perceptual reversals per second on the left and the corresponding ROC curve on the right. For each bin, the number of uninformed subjects is shown in pink on the left and the informed subjects in cyan on the right. Each point on the curve corresponds to a threshold value for the rate of perceptual reversals (color map), and the axes represent the proportion of uninformed subjects whose rate exceeds that threshold (x-axis), and the proportion of informed subjects whose rate exceeds that threshold (y-axis). Each point on the curve corresponds to a threshold value for the rate of perceptual reversals (color map), and the axes represent the proportion of uninformed subjects whose rate exceeds that threshold (x-axis), and the proportion of informed subjects whose rate exceeds that threshold (y-axis).

the sphere was ambiguous (these participants may have misunderstood the instruction or the survey question; see Discussion), and numerous participants in the uninformed condition reported thinking that it was (these participants may have been familiar with ambiguous stimuli already). We performed an ART (aligned rank transform) ANOVA (Wobbrock et al., 2011), a non-parametric equivalent of two-way ANOVA, with condition and responses to this question as factors. Results revealed that while the effect of condition persists [$F(1, 470) = 6.82$, $p = .009$], whether subjects thought the sphere was ambiguous showed no effect [$F(1, 470) = 2.26$, $p = .133$]. Nonetheless, to ensure that we do not overlook a potential influence of self-reported knowledge, depending on the combination of condition and response to question 6 of the post-experiment survey (See Appendix B), we further divided the participants into four sub-groups to investigate within-condition differences.

Subjects within the uninformed condition who reported 'yes' ($N = 135$) reported 35.71 reversals ($SD = 27.28$) at a rate of 0.20 reversals per second ($SD = .15$) on average while subjects in

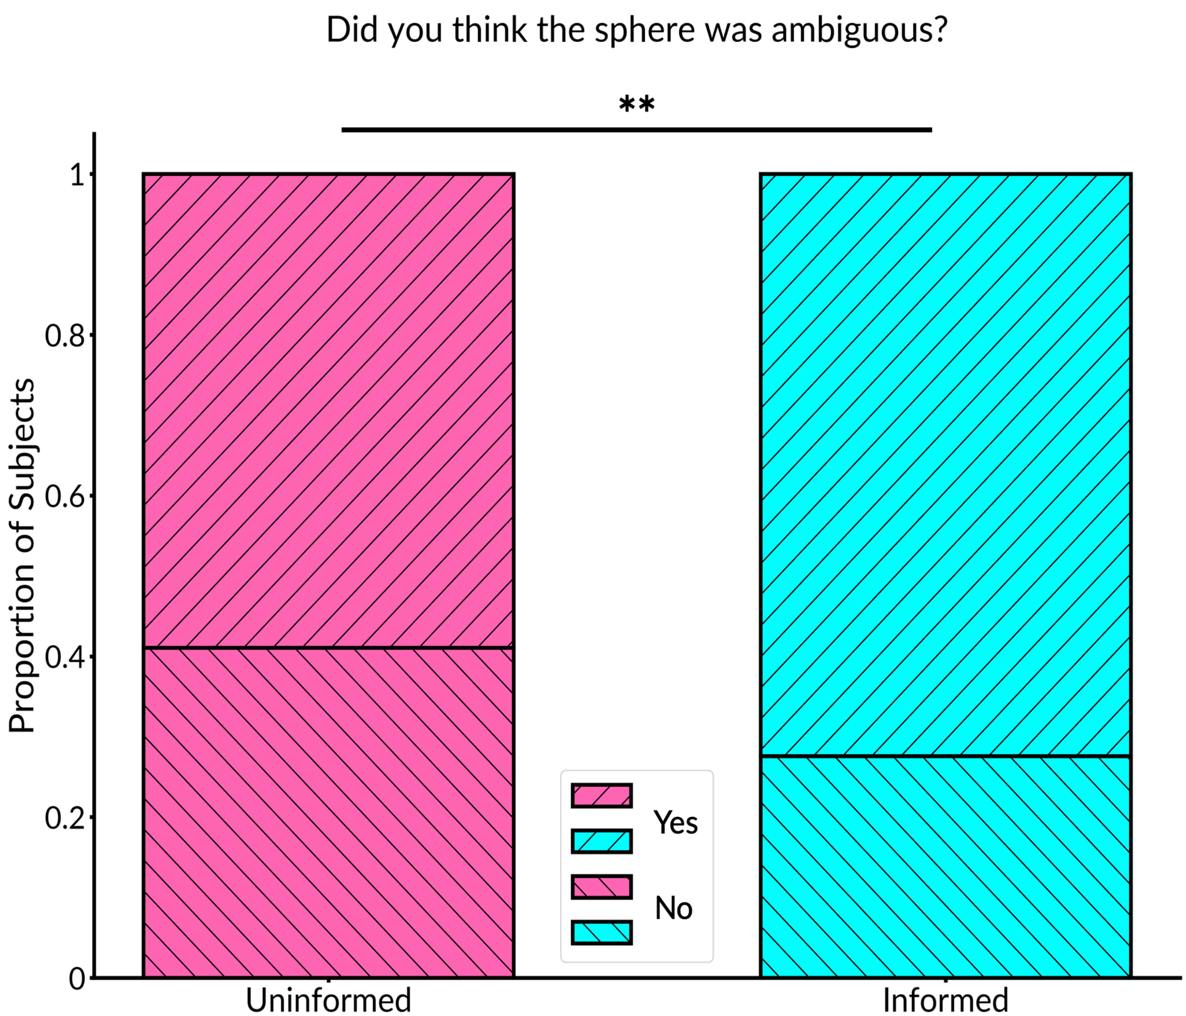

Figure 3.3: Participants' responses to question 6 of the post-experiment survey. A Fisher's Exact Test revealed that a larger proportion of subjects in the informed condition than in the uninformed condition. * $p < .05$, ** $p < .01$, *** $p < .001$

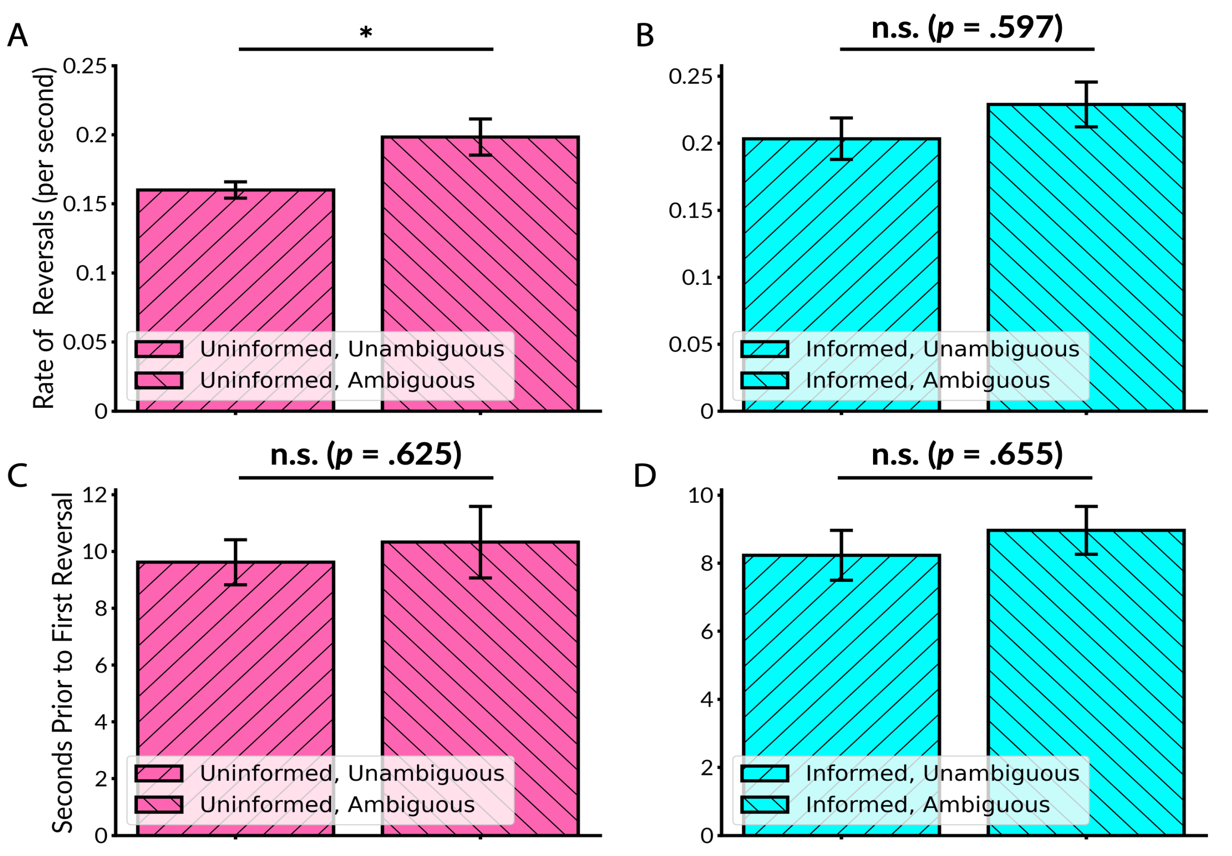

Figure 3.4: Within-condition results. The left bars represent the average rate of perceptual reversals per second. Subjects who thought the sphere was ambiguous reported significantly more perceptual reversals than those who did not. The right bars represent the average time in seconds taken prior to participants' initial reversal. No difference was found in this comparison. * $p < .05$, ** $p < .01$, *** $p < .001$

that condition who said 'no' (N = 94) reported 28.81 reversals (SD = 10.41) at a rate of 0.16 reversals per second (SD = .06). Among the uninformed subjects, we found this difference to be significant: more perceptual reversals among those who reported 'yes' than those who did not (Fig. 3.4A, p = .049). Similar to between-condition analysis, there was considerable overlap in across-participant distribution of reversal rate (Fig. 3.5A). We again performed a ROC analysis to examine the magnitude of this effect and found an AUC of 0.576 suggesting a relatively minor effect of knowledge (Fig. 8B).

For the informed condition, subjects who thought the sphere was ambiguous (N = 176) reported

41.21 reversals (*SD* = 39.97) at a rate of 0.23 reversals per second (*SD* = .22) on average while those who did not (N = 67) reported 36.61 reversals (*SD* = 22.76) at a rate of 0.20 reversals per second (*SD* = .13). In the informed condition this difference was not significant (Fig. 3.4B, $p = .597$).

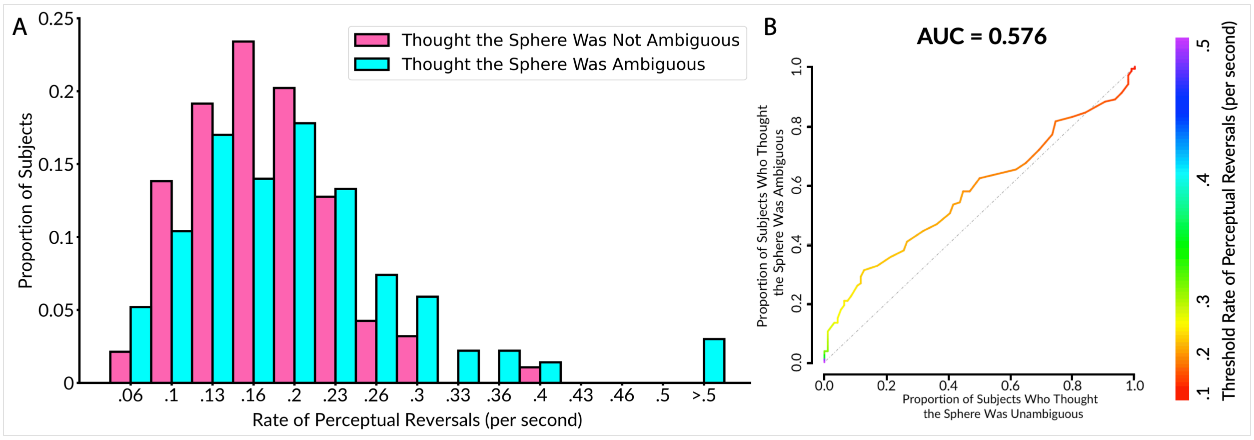

Figure 3.5: Within-condition results. The left bars represent the average rate of perceptual reversals per second. Subjects who thought the sphere was ambiguous reported significantly more perceptual reversals than those who did not. The right bars represent the average time in seconds taken prior to participants' initial reversal. No difference was found in this comparison. * $p < .05$, ** $p < .01$, *** $p < .001$

For the time it took participants to report their first perceptual reversal, we again conducted an ART ANOVA with condition and response to this question as the factors (Wobbrock et al., 2011). This revealed that neither condition [$F(1, 470) = 0.81$, $p = .368$], nor whether subjects thought the sphere was ambiguous showed an effect [$F(1, 470) = 0.84$, $p = .361$]. These results were confirmed by the following Mann Whitney U tests. We found no significant differences depending on whether they thought the sphere was ambiguous. On average, uninformed subjects who reported 'yes' took 10.33 seconds before reporting their first perceptual reversal (*S*D = 14.64) and those who reported 'no' took 9.62 seconds (*SD* = 7.72). Informed subjects who reported 'yes' took 8.97 seconds before reporting their first perceptual reversal (*SD* = 9.35) and those who reported 'no' took 8.23 seconds (*SD* = 6.02). The difference was not significant in either the uninformed (Fig. 3.4C, $p = .67$) or the informed condition (Fig. 3.4D, $p = .38$).

Participants' responses to the rest of the post-experiment survey were also examined, both in terms of response differences between conditions, and in terms of relations between those responses and reversal dynamics. In none of those cases were survey responses predictive of reversal rate or time taken to first reversal, but we did observe some surprising between-condition differences in terms of the proportions of ’yes’ versus ’no’ responses. Those differences will be addressed in our discussion section.

CHAPTER 4

DISCUSSION

The current study examined the effect of explicit knowledge on bistable perception of an ambiguous rotating sphere. Subjects observed the sphere and continuously reported their perceived direction of rotation. We measured observers' knowledge of bistable perception, and their experience with this stimulus, using a post-experiment survey, and we also manipulated this knowledge as a between-condition manipulation. We examined the relation between participants' knowledge about the reversibility of the stimulus and the dynamics of perceptual reversals reported while viewing the bistable stimulus. Our results suggest that knowledge has no impact on the time until the first reversal and has a small but significant effect on the rate of perceptual switches.

4.1 Modest effect of knowledge

Between-condition comparisons of perceptual reversals revealed significantly more reversals among the informed participants than their uninformed counterparts. Specifically, participants who knew that the sphere was ambiguous (i.e. were aware that their fluctuating perception of the ambiguous sphere was subjective) reported more reversals. Corroborating existing literature that argues for a facilitatory role of explicit knowledge, this finding suggests that knowledge aided subjects in switching between perceptual alternatives of the rotating sphere. However, this effect is much smaller compared to previous reports of between-condition differences. For instance, unlike uninformed subjects in prior work who reported few to no reversals (Girgus et al., 1977; Rock & Mitchener, 1992), subjects in the uninformed condition of the current study, even those whose survey responses confirmed that they were unaware of the ambiguity of the sphere or the subjective nature of their perceptual alternations, still experienced 28.8 reversals on average within the 180 seconds of our experiment at a rate of 0.16 reversals per second (Fig. 3.4A). A related point is that, in the past, learning the perceptual alternatives of the ambiguous figures led almost all observers to report multiple reversals, even if most of them had not reported any reversal at all prior to being informed (Girgus et al., 1977; Rock & Mitchener, 1992). In our data, on the other hand, between the uninformed and informed conditions, only a much more subtle shift is present in the overall sample toward the tendency to report more perceptual reversals (Fig. 3.2A). This is further supported by the ROC curve with an area under the curve of just 0.579 (Fig. 3.2B). The same trend is found again in our within-condition results using self-reported measures of knowledge. Within the informed condition, whether subjects thought the sphere was ambiguous at the time of their task was not significantly predictive of the number of perceptual reversals (Fig. 3.4B). Within the uninformed group this effect was significant, but still relatively modest as indicated by the ROC curve with an area under the curve of just 0.576 (Fig. 3.5B).

Another important aspect of the impact of knowledge that has been reported in prior studies but is not observed in the current study, is that knowledge also facilitates the first occurrence of perceptual reversal: in existing studies, those who were informed tended to reverse earlier following the start of the trial (Girgus et al., 1977; Rock & Mitchener, 1992). In our experiment, this would mean that informed observers should take less time than the uninformed to report their first perceptual reversal, but we did not find this (See Fig. 3.1B, Fig. 3.4C & D). In contrast to the majority of prior work emphasizing that informed subjects reversed 'quickly' (Girgus et al., 1977; Rock & Mitchener, 1992), our data suggests that it took similar amount of time for all subjects to report their first reversal, irrespective of whether they were informed regarding the sphere's ambiguity (Fig. 3.1B). The same trend persists in our within-condition results using self-reported measures of knowledge such that whether subjects thought the sphere was ambiguous is not predictive of time taken prior to first reversal in either condition (Fig 3.4C & D). In fact, almost all subjects (97.5%) in our experiment reported their initial perceptual reversal within 30 seconds of viewing the rotating sphere. Particularly contradicting prior work is that 223 out of 229 subjects in the uninformed group (97.4%) reversed within the first 30 seconds. This proportion for the uninformed subjects was approximately only half in previous studies that also used 30-second viewing periods, even in some cases when the reporting of perceptual changes was prompted (Girgus et al., 1977; Rock & Mitchener, 1992). Overall, our main results suggest a significant but modest relation between knowledge, either experimentally manipulated or measured through self-reported, and the dynamics of perceptual reversals during the viewing of the rotating sphere.

4.2 Discrepant role of knowledge between current and prior work

The differences between the current findings and existing literature could be attributed to a myriad of factors, chief among which is our decision to adopt the rotating sphere as stimulus. Historically, static ambiguous figures were popular choices for work on this particular topic. As outlined in the introduction, we chose the rotating sphere because it allowed us to use a task that was less likely to lead observers to wonder about hidden aspects of the stimulus or actively search for clues. Specifically, the rotating sphere allowed a task that subjects could perform, and could feel they were performing correctly, irrespective of whether they experienced perceptual reversals. However, the sphere does not fit into the typical categories of ambiguous figures used in prior work on this topic (Wimmer et al., 2011). These categories include content ambiguous (e.g., duck/rabbit, see Kosegarten & Kose, 2014), perspective ambiguous (e.g., the Necker cube in Girgus et al., 1977, or the Mach book in Rock & Mitchener, 1992), or figure-ground ambiguous (e.g., face/vase, see Kosegarten & Kose, 2014; see Wimmer & Doherty, 2007 for this classification of ambiguous figures). Out of those three categories, only the one related to perspective somewhat applies to our stimulus, but due to the rotation our stimulus is still quite different from perspective-ambiguous figures used in related work. It is possible that perception of the sphere, and of other ambiguous stimuli that involve three-dimensional rotation, involves different mechanisms than perception of static figures. Indeed, inter-observer correlation studies that center on the rate of perceptual reversals have shown reversal rates to correlate strongly among stimuli that do involve three-dimensional rotation, while correlations between such stimuli and stimuli without rotation are often absent or small (Brascamp et al., 2018; Cao et al., 2018; Jagtap & Brascamp, 2020). Such findings suggest that what applies to the rotating sphere may not be true for the stimuli used in existing work on the role of knowledge, thus potentially explaining the comparatively minor role of knowledge in our findings.

We will speculate about one possible explanation as to why perception of some stimuli used in previous work might involve such a large effect of knowledge. Specifically, perceptual interpretations for content-ambiguous figures may not always be convincing or salient if the figure is not presented in a certain fashion, in which case it helps to inform the participant about the interpretations. For example, several prior studies included a static ambiguous outline drawing called the chef/dog illusion Figure. D1, also see Rock, 1956; Rock et al., 1994; Rock & Mitchener, 1992, for more detail. While this figure is essentially content-ambiguous meaning that the outline can be seen as representing a chef or a dog, neither of the two perceptual interpretations may be compelling to observers without them performing some sort of mental rotation (Fig. D1 bottom). In support of this possibility, in order to reliably promote the perception of a particular interpretation of this figure (and of a previous iteration of the figure) authors of previous studies deliberately presented the figure in a particular orientation (Gibson & Robinson, 1935; Rock, 1956) . Also consistent with this possibility, Rock et al. (1994) reported that, out of uninformed participants who had previously been shown the illusion at a neutral orientation that did not favor either interpretation, only 14% acknowledged having seen the figure before when subsequently shown a variant at a different orientation (and color) designed to promote a specific percept. Such findings suggest that the intended perceptual interpretations of the chef/dog figure may not be compelling (especially without a specific orientation manipulation), which raises a possible explanation for the extensive role of knowledge that has been observed for that stimulus and potentially other content-ambiguous stimulus. In particular, for such stimuli knowledge may be required for the observer to actively seek out, or even acknowledge, either or both of the interpretations, whereas for our sphere stimulus the percepts may have a tendency to simply present themselves to in the mind's eye.

Another possible contributing factor to the discrepancy between our work and earlier findings could be the non-negligible amount of pre-existing knowledge about bistable perception in our participants. In the current study, all subjects were undergraduate students, which is consistent with most other studies (Kosegarten & Kose, 2014; Rock et al., 1994; Rock & Mitchener, 1992). Approximately 70% of our participants reported never having seen a similar sphere to the rotating sphere used here, 60% of them reported having heard of ambiguous displays prior to the experiment (Fig. D2), indicating that many of our participants had some knowledge about perceptual bistability. Even though extensive pre-existing knowledge is clearly suboptimal when it comes to our objective of experimentally manipulating participant knowledge, we do not think it forms a good explanation for the relative modesty of knowledge effects in our data. After all, we have shown that self-reported knowledge of bistability is not a better predictor of percept dynamics than experimentally induced knowledge (Figures 3.4 and 3.5). (See Appendix E for similar results based on other survey questions than the one used in Figures 3.4 and 3.5).

While not in line with some prior work with regard to the importance of knowledge in bistable perception, our findings are not at all inconsistent the existence of some influence of knowledge. Such an influence is not unexpected in light of other work on the influence of higher-order cognitive factors. For instance, many past studies have reported that attention (Chong et al., 2005; Meng & Tong, 2004), and intention (Liu et al., 2012; van Ee et al., 2005) can affect the dynamics of perceptual reversals. In the current study we did not instruct our subjects to direct attention to any specific part of the stimulus, nor to intentionally reverse or hold a percept. However, this does not exclude the possibility that knowing about the subjective nature of direction reversals motivated some subjects to do one of those things.

4.3 Self-reported experience and time to initial reversal

In our main text we addressed the only survey question for which participants' responses appeared to significantly predict perceptual reversal dynamics (Fig. 3.3). This was also the question that most directly indexed whether participants had thought the sphere was ambiguous during the experiment. We found no significant relations of that kind for any of the other survey questions. However, for one question we did find a between-condition difference in response proportions that surprised us. In particular, a larger proportion of subjects in the informed condition reported having seen a similar transparent rotating sphere prior to the experiment ($p < .001$, Fig. D2). This was unexpected because the experiment instructions clearly can have no influence on what participants had seen prior to the experiment. One possible explanation is that participants had pre-existing knowledge of other spheres similar to our stimulus and that our instructions for the informed condition may have confirmed their suspicion upon first impression of the sphere. This could, in theory, result in a larger proportion of the informed participants thinking they had seen a similar sphere prior to the experiment. Importantly, it is virtually impossible to preemptively assess participants' knowledge without somewhat revealing to them the goal of the study, just like it is virtually impossible to prevent the instructions from influencing their responses to the survey. Thus, even though asking these questions post-task may have led their responses to sway in a certain direction, we think that this is an important compromise to ensure the effectiveness of our manipulation. We will also reiterate that responses to this survey question, like those to all survey questions aside from the one discussed in our results section, did not significantly predict percept dynamics.

A second noteworthy data pattern related to our survey questions is that although our between-condition manipulation did significantly affect whether participants reported thinking the sphere was ambiguous during the experiment, responses were not cleanly divided between groups: many participants in the uninformed condition responded 'yes' to the question whether they knew the sphere was ambiguous, and many participants in the informed condition responded 'no' (See Fig. 3.3). This suggests that some of our participants in the uninformed condition had pre-existing knowledge and therefore realized that the sphere was ambiguous without being explicitly informed of it. As for the informed participants who thought the sphere was unambiguous (most of those participants also indicated they never seen a similar sphere; Fig. D2), it could mean that they misunderstood our instructions or the survey questions. In either case, these data patterns do not strongly affect our conclusions, considering that neither condition, nor the response to this question, strongly affected percept dynamics.

In our main results, cumulative time until initial perceptual reversal was defined starting at the beginning of the first real trial, but one could argue that participants had already been exposed to the stimulus prior to the trials, namely during the practice trial that was included in both conditions.

To ensure differences in time of first reversal are not overlooked, we also estimated another time to first reversal that includes the practice trial, as subjects may have reported their real initial reversal during practice. Specifically, this represents time in seconds from the beginning of the practice trial to each participant's first report of perceptual reversal, including reversals during practice. Consistent with the other measure of time to first reversal, between-condition comparison of this measure yields no significant difference (p = .485, Fig. D5).

CHAPTER 5

CONCLUSION

We were interested in testing the effect of conscious knowledge on the perception of ambiguous displays. We examined whether knowledge influences the occurrence of perceptual reversals when viewing a reversible rotating sphere. In line with previous work, our findings suggest that the rate at which perceptual reversals occur was affected by knowledge. However, in contrast with previous work, the impact of knowledge was much smaller in our data. The influence of knowledge, moreover, did not extend to the time taken prior to the initial reversal. In conclusion, the current findings, while partly supporting a role of knowledge in facilitating reversals in the perception of ambiguous visual input, do not confirm the observation in previous studies that the impact of knowledge extends beyond a relatively minor modulation of reversal rate, at least not for the stimulus we used. The fact that the impact is not zero either is consistent with the popular notion that both top-down and bottom-up mechanisms jointly drive perception during perceptual multistability.

APPENDICES

APPENDIX A

EXPERIMENTAL INSTRUCTIONS

A.1 Uninformed condition

1. (Displays Fig 2.1A) This image illustrates what you will see during this experiment: a sphere made of many small dots, that is rotating around its axis.

2. Your task will be to indicate the direction in which the sphere is rotating. In other words: the direction in which the side closest to you is moving.

A.2 Informed condition

1. (Displays rotating sphere animation) At the center of the screen, you see what appears to be a rotating sphere, correct? Take a second and think about this: is the sphere rotating to the left or the right side of the screen?

2. Whichever direction you were thinking, you were right. In reality, of course, your screen is flat, and the animation you are viewing is two-dimensional whereas a real sphere would be three-dimensional and impossible to appear on a flat screen. The reason you may be perceiving a three-dimensional sphere is because your mind recognizes that the animation is consistent with the motion a sphere goes through, even though the animation itself is two-dimensional.

3. But here is the trouble: the animation is consistent with a sphere that rotates to the left, but it is also consistent with a sphere that rotates to the right. So, when we asked which direction of motion you saw, either answer was correct! The animation is ambiguous. In fact, as you pay closer attention to the sphere, you will likely see that your answer changes because your perception fluctuates: sometimes you see it rotating in one direction, whereas at other times you see the other. But it is always the same animation. Your mind cannot decide,

so to speak, in which direction the sphere rotates, so it settles on different interpretations at different times.

4. And that is what your task will be about: you will be asked to report which direction you perceive the sphere to be rotating toward at different times. This means that you will have to press a different key each time your interpretation of the sphere changes or, in other words, each time the rotation direction you perceive abruptly reverses.

APPENDIX B

POST-EXPERIMENT SURVEY

1. Before you started the experiment, had you ever seen a transparent, dotted, rotating shape similar to this one before? (Yes/No)

2. (An upright-positioned animated schematic that abruptly reverses between the two unambiguous versions of the sphere is displayed at the center of the screen.) Here you see a schematic depiction of abrupt reversals in the direction in which the sphere rotates around its axis: the sphere is shown rotating clockwise at one moment, and counter-clockwise a split second later. During the experiment, did you ever notice abrupt direction reversals like these? (Yes/No)

3. (This question only appears if subject responded 'Yes' to the previous question.) If somebody told you that the abrupt direction reversals that you saw were not actually happening on the screen, but only in your head, would you believe it? (Yes/No)

4. (An animated schematic that gradually changes only its orientation is displayed at the center of the screen.) Here you see a schematic depiction of gradual changes in the orientation of the sphere, the sphere is shown with its axis leaning leftward at one moment, and rightward a little while later. During the experiment, did you ever notice gradual orientation changes like these? (This is a catch question: subjects in both conditions should have seen this, Yes/No)

5. Have you ever heard of ambiguous displays? That is, of images or animations that are ambiguous, so that one and the same image or animation can be perceived in two different ways? (Yes/No)

6. During the experiment, did you think that you were watching an ambiguous display? (Yes/No)

APPENDIX C

DESCRIPTIVE STATISTICS

Table C.1: Descriptive statistics arranged by conditions.

Descriptive Statistics	Condition	Total Number of Reversals	Rate of Reversal (per second)	Timing of the Initial Reversal (seconds)
N	Informed	243	243	243
	Uninformed	229	229	229
Mean	Informed	39.9	.22	8.76
	Uninformed	32.9	.18	10.0
Standard Deviation	Informed	36.1	.20	8.56
	Uninformed	22.2	.12	12.3
Minimum	Informed	8	.04	1.62
	Uninformed	1	.006	1.97
Maximum	Informed	314	1.74	68.8
	Uninformed	222	1.23	158

Table C.2: Descriptive statistics arranged by self-reported knowledge.

Descriptive Statistics	Q6	Total Number of Reversals	Rate of Reversal (per second)	Timing of the Initial Reversal (seconds)
N	Yes	311	311	311
	No	161	161	161
Mean	Yes	38.8	.22	9.56
	No	32.1	.18	9.04
Standard Deviation	Yes	35.1	.20	11.9
	No	17.1	.095	7.08
Minimum	Yes	1	.006	1.62
	No	4	.022	1.80
Maximum	Yes	314	1.74	158
	No	171	.95	46

APPENDIX D

SUPPLEMENTARY FIGURES

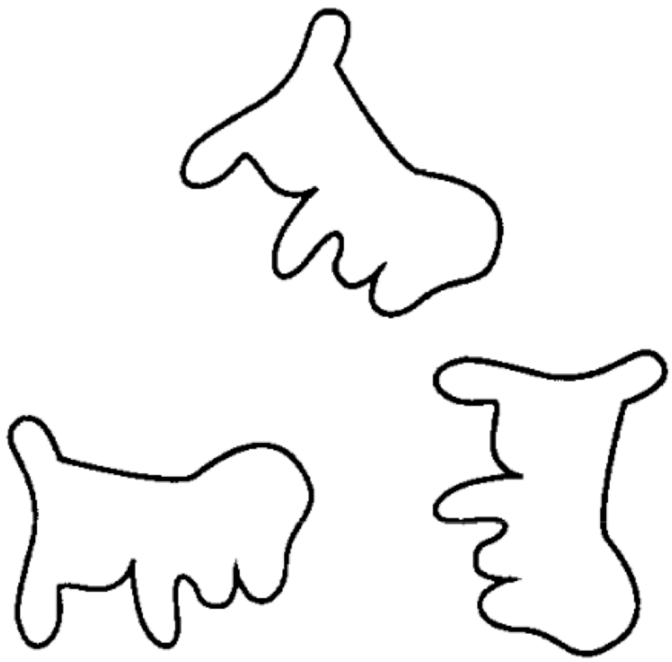

Figure D.1: The chef/dog illusion by Rock (1956). The top left figure shows the 'ambiguous' image that is typically presented to observers. The bottom left figures show the 'unambiguous' interpretations of the figure, although requiring some degree of mental rotation.

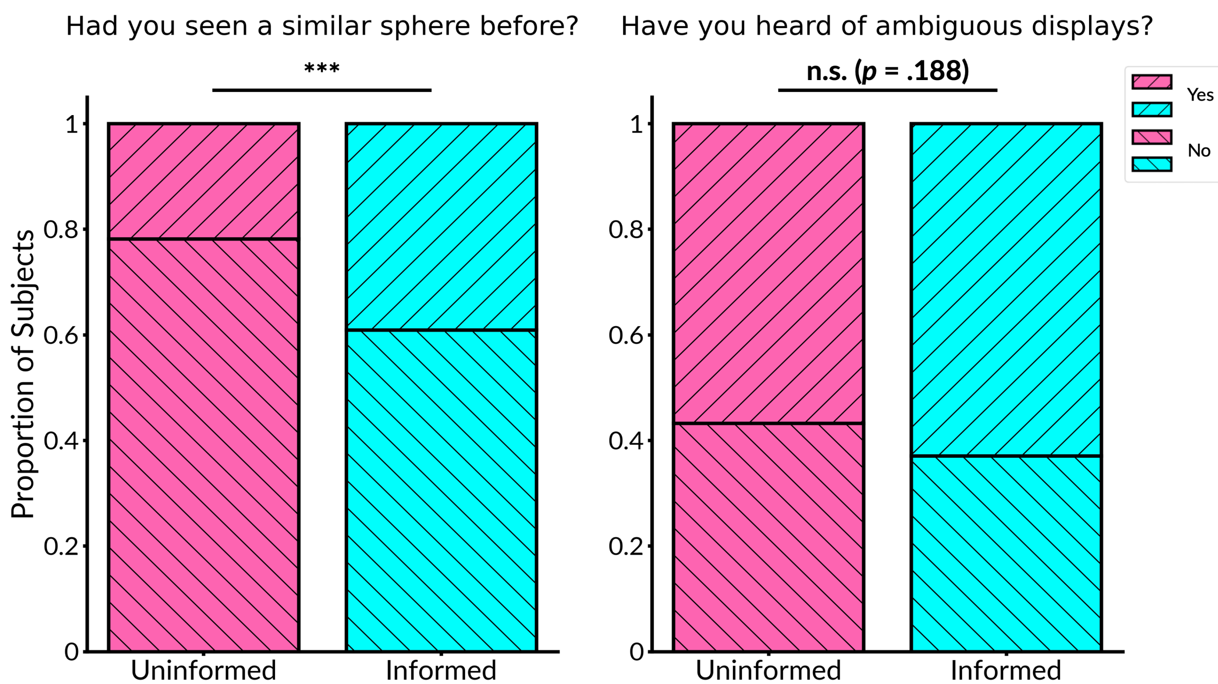

Figure D.2: Participants' responses to question 1 and question 5 of the post-experiment survey. A Fisher's Exact Test revealed that a larger proportion of subjects in the informed condition reported having seen a similar sphere prior to the experiment than in the uninformed condition. No difference was found for question 5. * $p < .05$, ** $p < .01$, *** $p < .001$

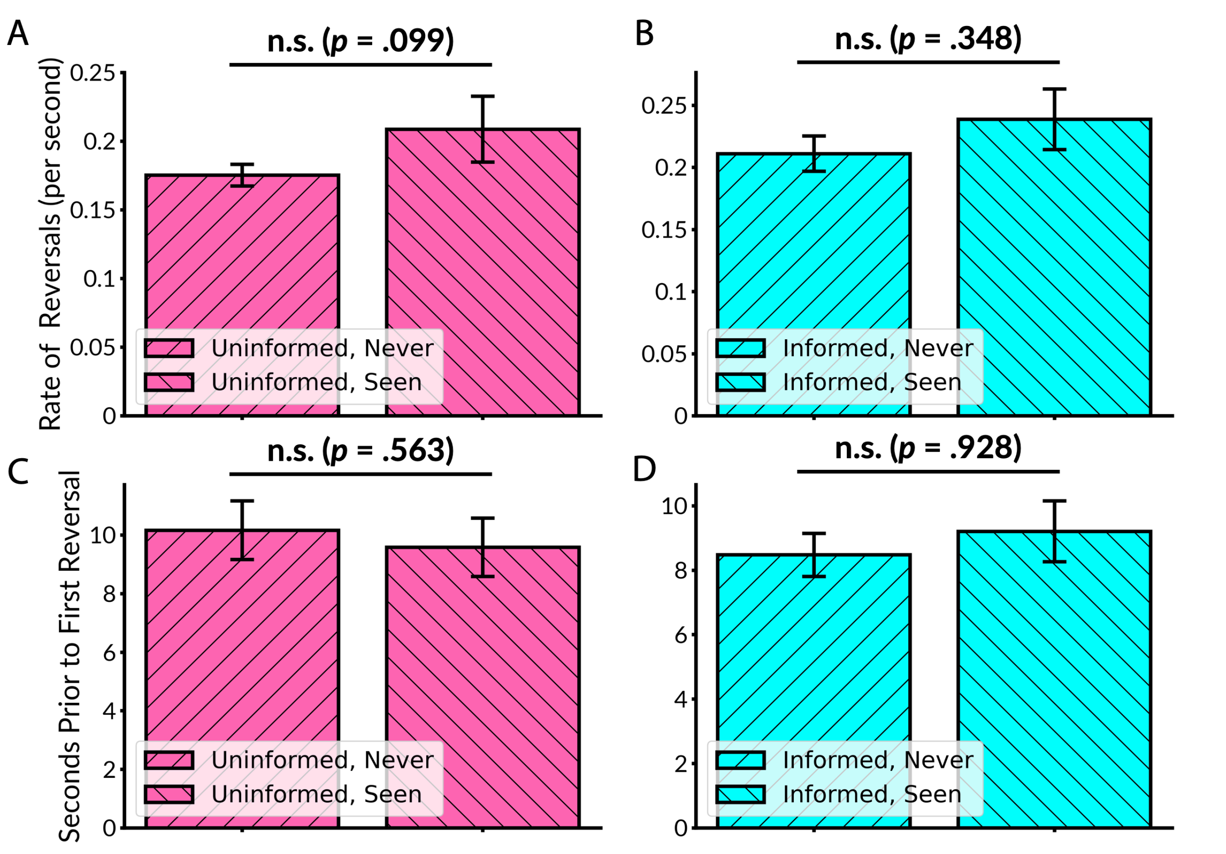

Figure D.3: Within-condition results arranged by response to question 1 of the survey, or whether participants had seen a similar sphere prior to the experiment. The left bars represent the average rate of perceptual reversals per second. The right bars represent the average time in seconds taken prior to participants' initial reversal. No difference was found. * $p < .05$, ** $p < .01$, *** $p < .001$

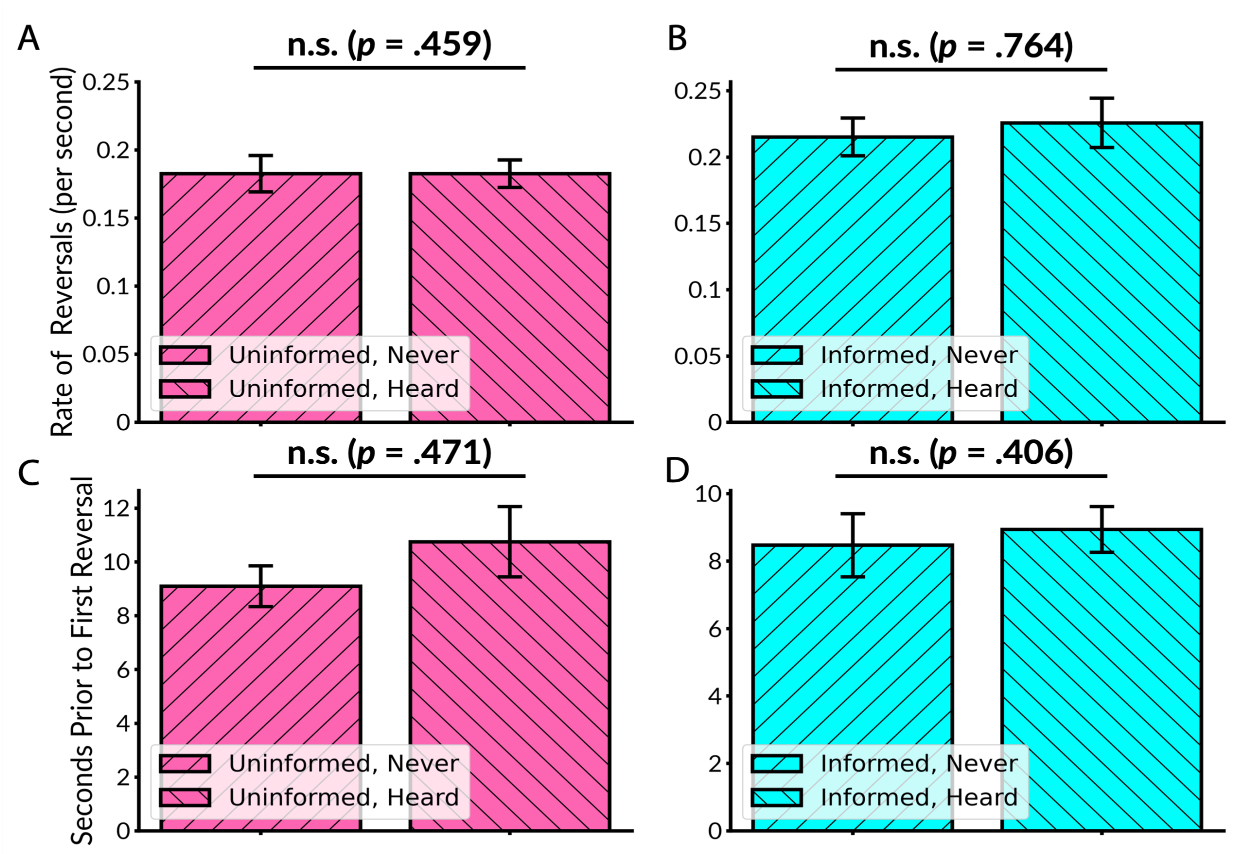

Figure D.4: Within-condition results arranged by response to question 5 of the survey, or whether participants had heard of ambiguous displays prior to the experiment. The left bars represent the average rate of perceptual reversals per second. The right bars represent the average time in seconds taken prior to participants' initial reversal. No difference was found. No difference was found. * $p < .05$, ** $p < .01$, *** $p < .001$

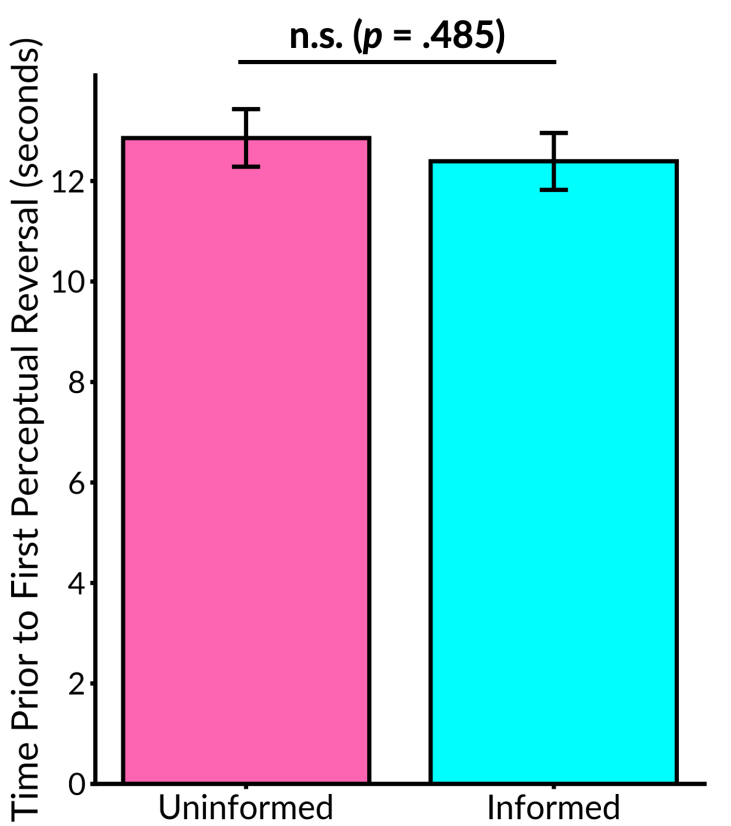

Figure D.5: Average time to first reversal including practice trial. This means that the cumulative time until first reversal for each participant includes practice time and reversal reported during practice. * $p < .05$, ** $p < .01$, *** $p < .001$

BIBLIOGRAPHY

BIBLIOGRAPHY

Ach, N. (1910). Uber den willensakt und das temperament: Eine experimentelle untersuchung.

Blake, R., Sobel, K. V., & Gilroy, L. A. (2003). Visual motion retards alternations between conflicting perceptual interpretations. *Neuron*, *39*(5), 869–878.

Brascamp, J. W. (2021). Controlling the spatial dimensions of visual stimuli in online experiments. *Journal of Vision*, *21*(8), 19–19.

Brascamp, J. W., Becker, M. W., & Hambrick, D. Z. (2018). Revisiting individual differences in the time course of binocular rivalry. *Journal of Vision*, *18*(7), 3–3.

Cao, T., Wang, L., Sun, Z., Engel, S. A., & He, S. (2018). The independent and shared mechanisms of intrinsic brain dynamics: Insights from bistable perception. *Frontiers in Psychology*, 589.

Chen, X., & He, S. (2004). Local factors determine the stabilization of monocular ambiguous and binocular rivalry stimuli. *Current Biology*, *14*(11), 1013–1017.

Chong, S. C., Tadin, D., & Blake, R. (2005). Endogenous attention prolongs dominance durations in binocular rivalry. *Journal of vision*, *5*(11), 6–6.

Gibson, J. J., & Robinson, D. (1935). Orientation in visual perception; the recognition of familiar plane forms in differing orientations. *Psychological Monographs*, *46*(6), 39.

Girgus, J. J., Rock, I., & Egatz, R. (1977). The effect of knowledge of reversibility on the reversibility of ambiguous figures. *Perception & Psychophysics*, *22*(6), 550–556.

Helmholtz, H. L. (1962). *Helmholtz's treatise on physiological optics*. Dover.

Hering, E. (1964). Outlines of a theory of the light sense.

Jagtap, A. R., & Brascamp, J. W. (2020). The role of 3d rotation in perceptual bistability. *Journal of Vision*, *20*(11), 937–937.

Kosegarten, J., & Kose, G. (2014). Seeing reversals in ambiguous images: To know or not to know? *Perceptual and Motor Skills*, *119*(1), 228–236.

Leopold, D. A., & Logothetis, N. K. (1999). Multistable phenomena: Changing views in perception. *Trends in cognitive sciences*, *3*(7), 254–264.

Lewin, K. (1922). *Der begriff der genese in physik, biologie und entwicklungsgeschichte*. Springer.

Li, Q., Joo, S. J., Yeatman, J. D., & Reinecke, K. (2020). Controlling for participants' viewing distance in large-scale, psychophysical online experiments using a virtual chinrest. *Scientific reports*, *10*(1), 1–11.

Liu, C.-H., Tzeng, O. J., 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.

McKnight, P. E., & Najab, J. (2010). Mann-whitney u test. *The Corsini encyclopedia of psychology*, 1–1.

Meng, M., & Tong, F. (2004). Can attention selectively bias bistable perception? differences between binocular rivalry and ambiguous figures. *Journal of vision*, *4*(7), 2–2.

Pastukhov, A., Lissner, A., & Braun, J. (2014). Perceptual adaptation to structure-from-motion depends on the size of adaptor and probe objects, but not on the similarity of their shapes. *Attention, Perception, & Psychophysics*, *76*(2), 473–488.

Rock, I. (1956). The orientation of forms on the retina and in the environment. *The American journal of psychology*, *69*(4), 513–528.

Rock, I., Hall, S., & Davis, J. (1994). Why do ambiguous figures reverse? *Acta psychologica*, *87*(1), 33–59.

Rock, I., & Mitchener, K. (1992). Further evidence of failure of reversal of ambiguous figures by uninformed subjects. *Perception*, *21*(1), 39–45.

van Ee, R., Van Dam, L., & Brouwer, G. (2005). Voluntary control and the dynamics of perceptual bi-stability. *Vision research*, *45*(1), 41–55.

Wimmer, M. C., Doherty, M. J., & Collins, W. A. (2011). The development of ambiguous figure perception. *Monographs of the Society for Research in Child Development*, i–130.

Wobbrock, J. O., Findlater, L., Gergle, D., & Higgins, J. J. (2011). The aligned rank transform for nonparametric factorial analyses using only anova procedures. *Proceedings of the SIGCHI conference on human factors in computing systems*, 143–146.