



4-2005

Seeing Beyond the Illusion: Dispelling the Grand Illusion Hypothesis

Melissa A. Ebbers

Follow this and additional works at: https://scholarworks.wmich.edu/masters_theses



Part of the Philosophy Commons

Recommended Citation

Ebbers, Melissa A., "Seeing Beyond the Illusion: Dispelling the Grand Illusion Hypothesis" (2005). *Master's Theses*. 1432.

https://scholarworks.wmich.edu/masters_theses/1432

This Masters Thesis-Open Access is brought to you for free and open access by the Graduate College at ScholarWorks at WMU. It has been accepted for inclusion in Master's Theses by an authorized administrator of ScholarWorks at WMU. For more information, please contact wmu-scholarworks@wmich.edu.



SEEING BEYOND THE ILLUSION:
DISPELLING THE GRAND
ILLUSION HYPOTHESIS

by

Melissa A. Ebbers

A Thesis
Submitted to the
Faculty of The Graduate College
in partial fulfillment of the
requirements for the
Degree of Master of Arts
Department of Philosophy

Western Michigan University
Kalamazoo, Michigan
April 2005

SEEING BEYOND THE ILLUSION:
DISPELLING THE GRAND
ILLUSION HYPOTHESIS

Melissa A. Ebbers, M.A.

Western Michigan University, 2005

The Grand Illusion hypothesis is a new form of skepticism about the nature of our visual experience: it *seems* to us as though our vision is everywhere detailed, distinct, continuous, in color, even “photograph-like,” but it is not. This position is motivated by developments in perceptual research, which have revealed new information about the functional structure of the visual system as well as the attention-dependent nature of perception.

My project is primarily deconstructive. I argue that the Grand Illusion hypothesis rests on problematic assumptions (motivated by the results from the relevant perceptual research), which ultimately leads to an incoherent formulation of what we ought to expect our vision to be like. By challenging these assumptions, our expectations are altered in such a way that predicts that we would have exactly the visual experiences that we do. Thus, we are not subject to the Grand Illusion.

ACKNOWLEDGMENTS

First and foremost, I would like to express my gratitude to the members of my thesis committee: Dr. John Dilworth, Dr. Timothy McGrew, Dr. Michael Pritchard, and Dr. Marc Alspector-Kelly. Their involvement in this process and genuine enthusiasm for the project contributed greatly to my success.

Secondly, I would like to thank my wonderful friends—including my classmates—and family (especially my mother and father) for their love and support during this process. Thank you for all of the fun and interesting discussions, as well as your kind words of encouragement.

Lastly, I would like to express great appreciation for the hard work and dedication of my advisor, Dr. John Dilworth—he certainly deserves special recognition for his contribution to this effort. His guidance and support, which included thorough comments, frequent meetings, advice, encouragement and helping me to keep a sense of humor about the project, has made this a highly rewarding experience. I am truly fortunate to have such a wonderful mentor.

Melissa A. Ebbers

Copyright by
Melissa A. Ebbers
2005

TABLE OF CONTENTS

ACKNOWLEDGMENTS	ii
CHAPTER	
I. AN OVERVIEW OF THE GRAND ILLUSION HYPOTHESIS	1
1.1 Introduction	1
1.2 The Grand Illusion Hypothesis: A New Skeptical Problem?	5
1.3 How Do We Determine Whether the Ordinary Perceiver is Subject to the Grand Illusion?	5
1.4 The Role of the Ordinary Perceiver.....	11
1.5 The Ordinary Perceiver’s Belief.....	12
1.6 What Does (and Does Not) Surprise the Ordinary Perceiver	12
II. THE PREDOMINATE METHOD WITHIN THE DEBATE FOR DETERMINING WHETHER THE ORDINARY PERCEIVER IS SUBJECT TO THE GRAND ILLUSION	16
2.1 Overview	16
2.2 Distinction Regarding Methods of Belief-Acquisition.....	18
2.3 The Method of Acquiring Experience-Based Beliefs.....	21
2.4 Other Candidates for the Belief.....	24
2.5 Further Consideration of Inference-Based Beliefs	26
2.6 Why Dennett Cannot Argue That the Ordinary Perceiver is Subject to the Grand Illusion.....	31
III. REEVALUATING THE ASSUMPTIONS THAT MOTIVATE THE GRAND ILLUSION HYPOTHESIS	38
3.1 What Assumptions Motivate the Hypothesis?	38

List of Figures—continued

3.2 A Distinction Between Kinds of Functional Flaws	43
3.3 Formulating the Predicted Nature of Experience	47
3.4 Alternatives to the Incoherent Formulation of the Predicted Nature of Experience	50
IV.A NEW FORMULATION OF THE PREDICTED NATURE OF EXPERIENCE	53
4.1 Considering the Affect that Each Kind of Flaw Has on Our Experience of the Other Kind of Flaw	53
4.2 Explaining the Surprise Within the New Account	55
4.3 Concluding Remarks	57
BIBLIOGRAPHY	59

1. AN OVERVIEW OF THE GRAND ILLUSION HYPOTHESIS

1.1 INTRODUCTION

The Grand Illusion is an illusion of which the layperson with normal vision (let us call her the Ordinary Perceiver) purportedly suffers. As characterized by Alva Noë, the Grand Illusion is a new kind of skepticism regarding how our visual experience seems to us. The advocates of this illusion argue that it seems to the Ordinary Perceiver that her vision is everywhere detailed, distinct, continuous, in color, photograph-like, etc.,¹ when her vision is, in at least some regions of her visual field, in fact poor, not detailed and even colorless!² As a means of determining whether such a disparity exists, discussion has focused around how to interpret the surprise—and lack of surprise—exhibited³ by the Ordinary Perceiver upon experiencing her functional flaws (blind-spot, change-blindness, inattentional-blindness, saccades, etc.) as well as during normal perception.⁴ This surprise is taken to be an indication that a belief is overthrown—that the Ordinary Perceiver was expecting something else—when she discovers her functional flaws. In what follows,

¹ I will also refer to this list of attributes as “photograph-like” (as an abbreviation of the list) This list of attributes is intended to be similar to the “snap-shot” conception addressed by Noë (*Is the Visual World a Grand Illusion?*, p.2)

² This assumption, that the predicted nature of vision is not everywhere detailed, distinct, continuous, in color, photograph-like, etc., is based on a further assumption that experiencing the functional flaws is the what we ought to expect that our visual experience would be like. I challenge this assumption in Chapters 3 and 4.

³ As discussed in §1.3, I am only arguing against the predominant method (in the literature on the debate) of determining whether the Ordinary Perceiver is subject to the Grand Illusion—through interpretation of the significance of the surprise, and lack thereof, exhibited by the her under difference circumstances that are relevant to the debate.

⁴ Noë argues that it is also important to note that the Ordinary Perceiver is *not* surprised by some of the adjustments (putting on her glasses, relocating, squinting, looking around, etc.) she must make to in order to perceive the desired information during normal perception. I address this in greater detail below.

I argue that the Grand Illusion hypothesis is motivated by expectations of what our visual experience must be like which are formed without consideration of sufficient relevant factors. Since these expectations are not met during our normal⁵ perceptual experience, as taken to be particularly evident given the surprise exhibited by the Ordinary Perceiver when she experiences any of her functional flaws, we assume that we are subject to an illusion. However, the surprise alone is insufficient to settle the matter since more than one belief could cause the surprise. To illustrate this point, let us discuss what motivates us to call the Ordinary Perceiver's experience illusory, namely expectations derived from recent research of the flaws.

Since the Grand Illusion hypothesis rests on expectations of what our perception must be like which are developed without considering all of the relevant factors, it is useful to attempt to formulate the predicted⁶ nature of perception based on what the Grand Illusion hypothesis takes to be illusory about the Ordinary Perceiver's perceptual experience as a means of drawing attention to the problematic assumptions which generate these expectations. My account, then, is two-fold: the first half is dedicated to what can be said regarding how the Ordinary Perceiver's visual experience seems to her, and the second discusses the implications of arguing that the Ordinary Perceiver's experience, as formulated by the Grand Illusion hypothesis, is illusory, concluding that the Grand Illusion hypothesis is motivated by unwarranted expectations (based on considering too little of the relevant data discovered through perceptual research), and arguing that once we generate our expectations by considering the affect of each of the kinds of flaw on our experience of the other, we can then see that we are not subject to an illusion—in fact, what we

⁵ By 'normal,' I mean how our experience seems to us during most, if not all, of our experiences—all but those occasions under which we perceive our structural flaws (which we do by means of non-standard usage of our perception)

⁶ See discussion of this term in §1.3.

experience is just what we would expect to experience, once we consider enough of the relevant (available) facts.

My inquiry during the first half of my project primarily addresses how the Ordinary Perceiver's vision seems to be, and incorporates other aspects of the debate only insofar as they aid in this project. I introduce the skeptical problem in Chapter 1 by addressing the functional flaws (§1.3) and what does (and does not) surprise the Ordinary Perceiver (§1.6). Through my analysis of how the Ordinary Perceiver's vision seems to her in Chapters 1 and 2, I discuss at length the role of the Ordinary Perceiver (§1.4), specifying what belief the Ordinary Perceiver must have in order to be subject to the Grand Illusion (§1.5-2.3) and whether there are any other candidates for the belief that causes the surprise (§2.4-2.5). I conclude the first half, by arguing that the case cannot be made for the alleged disparity between how the Ordinary Perceiver's vision seems to her (everywhere detailed, distinct, continuous, in color, photograph-like, etc.) and the predicted nature of her vision (discontinuous, colorless and without detail in some regions, etc.), and, as such, a new account is needed to explain the Ordinary Perceiver's reaction to various conditions. Since we cannot develop an account to explain this disparity of experience, in Chapter 3 I suggest that a new approach is adopted by the debate—re-examining our assumptions in an attempt to discover whether we may alter our expectations in such a way as to eliminate the Grand Illusion hypothesis in favor of a new explanation of our perceptual experience.

The second half of my project, then, is dedicated to what the Grand Illusion hypothesis implies about the predicted nature of perception (§3.1), which, if I am right, renders the set of expectations (generated by the assumptions made about the results of the research on the functional flaws) incoherent, illuminating an internal

problem with the formulation of the hypothesis (§3.3). In order to draw out the incoherence, I make a distinction between two kinds of functional flaws based on how they are generated (§3.2). Since it is claimed that the surprise exhibited by the Ordinary Perceiver when she experiences each kind of flaw (which are distinct kinds since they are generated by separate aspects of the perceptual system), then, since they are argued separately to be indicative of the Grand Illusion hypothesis, they admit of a distinct formulation of both how they seem to the Ordinary Perceiver, as well as of the predicted nature of perception. In other words, once the distinction is made between the two kinds of flaws, we can evaluate what the claim that the Ordinary Perceiver is having an illusory experience implies for the formulation of the predicted nature of perception regarding each kind of flaw. The two formulations of the predicted nature of perception are compared in an attempt to integrate them into one unified formulation of the predicted nature of perception. However, I argue that, once we attend to both kinds of functional flaws, the formulation of the predicted nature of perception becomes incoherent (§3.3).

Though my account is primarily deconstructive, I conclude by arguing that direction can be found for a new formulation of our expectations and the predicted nature of perception (§4.1). Once we consider both kinds of flaws (and the effect that they will have on our overall perceptual experience), we are able to formulate one account of the predicted nature of perception (rather than two separate formulations—one for each kind of flaw, and is based solely on the expectations that we would generate by virtue of that which we know about that flaw independently of how any other feature(s) of perception might effect it—which cannot be integrated into a single account). The expectations generated by the new formulation of the predicted nature of perception are such that we need no longer entertain the possibility that we are

subject to an illusion—this account entails that we should not expect to experience our flaws during normal perception (§4.3).

1.2 THE GRAND ILLUSION: A NEW SKEPTICAL PROBLEM?

Noë argues that the Grand Illusion hypothesis is a new skeptical problem: skepticism about our perceptual experience.⁷ The worry is whether our perceptual experience is different than it seems to us. The alleged disparity is that our vision seems to us to be everywhere detailed, distinct, continuous, in color, photograph-like, etc., but this experience does not accord with the predicted nature of vision, which is not everywhere detailed, distinct, continuous, in color, photograph-like, etc., due to the impairing effects of the functional flaws on our visual system. This skepticism is motivated by the surprise exhibited by the Ordinary Perceiver when she discovers one of her functional flaws. Since she is surprised, we infer that a belief was overturned—a belief, the contents of which lead her to expect that her vision would seem to her as it always had: detailed, distinct, continuous, in color, photograph-like, etc. In other words, the surprise at discovering one of her functional flaws indicates that her belief about the photograph-like nature of her vision was overturned at that time. Thus, we worry whether we ought to be skeptical about our visual experiences, given the alleged disparity between the predicted nature of our visual experiences and how they seem to us.

1.3 HOW DO WE DETERMINE WHETHER THE ORDINARY PERCEIVER IS SUBJECT TO THE GRAND ILLUSION?

Throughout the present critique of the Grand Illusion hypothesis, I restrict the scope of my evaluation to the predominate method within the debate for determining

⁷ Noë, *Is the Visual World a Grand Illusion?*

whether the Ordinary Perceiver is subject to the Grand Illusion.⁸ As such, I take issue with using the surprise exhibited by the Ordinary Perceiver when she first discovers her flaws as indication that she is subject to the Grand Illusion. In order to analyze the Grand Illusion hypothesis in the proposed fashion, we must first introduce some terms we will be using to discuss the debate. Those terms are used to must lay out the problem. Once we have clearly formulated the problem, we will be in a position to critically review the it.

Since it is my contention that the Grand Illusion hypothesis is motivated by a disparity between our expectations regarding what we would expect (based on the recent results from perceptual research), I will define my terms to reflect this. I propose the term ‘predicted nature of experience’ to refer to what we would expect our experiences to be like (relative to some set of data from which we form our expectations). Whatever assumptions regarding what we expect that our experience would be like (that we make based on the data that we are considering) are collectively considered to be the predicted nature of experience.

In contrast, an illusion (as the term is, arguably, used within the debate) is any experience which deviates from the respective predicted nature of experience. Though this is a non-standard usage of the term, something like this appears to be the implicit use of illusion within this debate. For example, given the data regarding the functional structure of the visual system (blind spot, photoreceptor distribution, etc.), we expect that our experience would be missing visual information in the region of the blind spot and would noticeably lessen in detail and color towards the periphery; these expectations comprise the predicted nature of experience. By virtue of the apparent failure of our normal experience to conform to our expectations, such

⁸ Thus, I am not speaking to other ways of arguing for the problem. I am only concerned with the general way in which the hypothesis is taken to be established.

experience is taken to be illusory.

The **functional flaws**⁹ of our visual system are aspects of the functional structure of our eyes that lead us to expect that our vision would not be everywhere detailed, distinct, continuous, in color, photograph-like, etc. The assumption appears to be that, given the existence of these flaws, we would expect that our vision is experienced as only having rich color and detail in the foveal¹⁰ region, lessening in color and detail towards the periphery, rapid sequences of snapshots, etc. Despite the presence of these flaws as part of the functional structure of our eyes, we do not notice them during normal perception¹¹ and are surprised to discover them.¹²

I am only addressing some of the flaws commonly found in the literature on the Grand Illusion that are relevant to my discussion of the disparity between how our experiences seem to us and the predicted nature of perception, such that only those flaws which can be experienced are to be taken into account of this debate. Saccades (a consequence of the function of photoreceptors), for instance, are discussed in the literature, but since I am interested in explaining the surprise exhibited by the Ordinary Perceiver when she experiences her flaws (and I am unaware of circumstances under which she could experience her saccades), I will not be discussing them.

One of the most popular of these functional flaws, the blind-spot, is the source of much excitement (and the cause of much speculation) for the visual theorist. The

⁹ I owe this term to discussions with John Dilworth. It is important to note that I am using the term “flaw” to connote that, in the debate, it is implied that these “flaws” are that which prevents us from having (the non-illusory) experience of vision as everywhere detailed, distinct, continuous, in color, photograph-like, etc.

¹⁰ The fovea of the eye is the region which we generally use during normal vision (see footnote 8). It contains a high number of cones (photo-receptors that enable us to have color vision), which allows us strong visual discriminatory power (good visual acuity, color, etc).

¹¹ By ‘normal perception,’ I mean one’s vision when one is not attempting to experience one of the flaws—this is the kind of vision we experience most, if not all, of the time.

¹² Under non-normal perceptual circumstances, such as experimental conditions.

blind-spot, which is caused by a region that contains no photoreceptors where the optic nerve connects to the eye and is located just 6 degrees from the center of the visual field (covering 16 degrees of the visual field), can be observed under certain circumstances. These circumstances differ from normal perception: during normal perception, we attend to that which is straight ahead of where our eyes are directed, while we must attend to that which is outside of the center of our vision in order to experience our blind-spot. Whatever objects or events that fall within our blind-spot are experienced (when we are *attending* to that region of our vision) as occluded from view. One of the reasons that the blind-spot is so interesting is that, though it is not far from the center of our visual field, we fail to experience it during normal perception.

Photoreceptor density and distribution is another functional flaw. We are virtually colorblind outside of the foveal region, and the level of perceptual detail diminishes from borders of the foveal region to the periphery. We do not notice either of these flaws during normal perception (as noted earlier, our experience seems to be highly detailed and richly colored during normal perception).

Change-blindness and inattentional-blindness are two related functional flaws. Both involve the failure to experience that to which we are not attending. In an experiment on inattentional-blindness, 'Gorillas in our Mist' (Simons, D. and Chabris, C., *Perception*, 1999, **28**), subjects were asked to watch a video recording of a basketball game, attending to the number of passes made by one team. During the game, a person in a gorilla costume walks onto the court, pauses to beat his chest, then exits the court. Though many of the participants correctly counted the number of passes, fewer than half noticed the gorilla. The participants were shocked when they watched the video again, observing the gorilla which they had previously not noticed.

Surprising results are also obtained through research on change-blindness.

Subjects fail to notice changes in stimuli that occur immediately in front of them. One such example is stimuli with “mud-splashes”: the subjects observe a picture, which changes every time a “mud-splash” (a shape that occludes part of the picture briefly) flashes.¹³ Subjects are surprised to learn that they failed to notice changes that occurred in the picture as they were looking at it. This sense of surprise—common to the discovery of change-blindness, inattentional-blindness and the blind-spot (as well as others)—acts as a basic step in arguing for the Grand Illusion.

Now that we have discussed the terms, we are in a position to lay out the problem. Recall that I argue that the Grand Illusion hypothesis is motivated by a deviations from the predicted nature of experience (as formulated based on the results from the research on the functional flaws). Thus, in order to determine whether she is subject to the Grand Illusion, we must answer the following questions:

- (1) Is there a discrepancy between the predicted nature of our experience and how our experience seems to us?
- (2) What is the predicted nature of our perceptual experience?
- (3) How does our perceptual experience *seem* to us?

The main question of the debate is (1). If the answer to (1) is no, then the Grand Illusion hypothesis (at least insofar as it accords with our formulation) is false; otherwise, further analysis is warranted. The answer to (1) will become more apparent in light of the answers to (2) and (3). Question (2) asks about the predicted nature of experience. According to the Grand Illusion, our vision, due to its functional flaws, is not everywhere photograph-like. In answering (2), we must address what facts about us (and the world) determine the predicted nature of vision, such that we are justified in asserting that our expectations regarding the predicted nature of perception are such

¹³ J. Kevin O'Regan offers an excellent example of change-blindness experiments that you can try for yourself at his website (<http://nivea.psycho.univ-paris5.fr/>).

that, when contrasted against the Ordinary Perceiver's experience (allegedly, that is it photograph-like), we are warranted in saying that she is subject to an illusion. Question (3) asks how experience really seems to the Ordinary Perceiver. The Grand Illusion hypothesis alleges that vision seems to the Ordinary Perceiver as though it is everywhere photograph-like. For present purposes, we must inquire how vision seems to the Ordinary Perceiver—specifically, whether it seems to her as though her vision is everywhere photograph-like. We may discover that there is no disparity between how our experience seems to us and the predicted nature of vision, in which case (1) is false. The only circumstances under which perceivers are subject to the Grand Illusion are if the answers to (2) and (3) pick out unique experiences/events/states/phenomena.¹⁴ If this is the case (if (1) is true), then we are justified in our skepticism at least to the extent that further analysis is reasonable. However, one further condition must be satisfied. Not only must there be a discrepancy between the answers to (2) and (3), but it must be the discrepancy specified by the Grand Illusion: it must seem to perceivers as though their vision is everywhere photograph-like, even though this is not the case.

Since we are concerned as to whether the Ordinary Perceiver is subject to the Grand Illusion, we need only inquire as to whether her vision seems to her to be everywhere photograph-like. If her vision does seem to her to be this way, then we are at least warranted in further investigation regarding the Grand Illusion.¹⁵ If it does not seem this way to her, then we have answered our question (at least regarding the Grand Illusion as formulated in this paper—which captures the skepticism discussed by Noë)—the Ordinary Perceiver does not suffer from the Grand Illusion.

¹⁴ I am using neutral language here to be inclusive to different theories of mind.

¹⁵ We would still need to determine whether the predicted nature of nature of experience is as it is claimed to be (in the formulation of the Grand Illusion).

1.4 THE ROLE OF THE ORDINARY PERCEIVER

One might wonder what are the conditions under which an individual may be counted as an Ordinary Perceiver, and why she is important to the debate. Surely her perceptual capacity is not so different from our own—what warrants preferential treatment for *her* beliefs? The Ordinary Perceiver is an individual who is not acquainted with visual theory. She is not aware that the structure of her visual system is such that it includes the functional flaws (mentioned above), and she is not part of the community that theorizes¹⁶ about perception. Dennett argues¹⁷ that she possesses “pre-theoretical” views about the nature of perception, based solely on how her experience seems to her. Let us grant for the sake of argument that those who argue as such are right—that she does have such commitments. These pre-theoretical views would be valuable, presumably, because they would not be biased by knowledge of visual theory, thus serving as an unbiased means for determining how experience “really” seems. Her role, then, is to shed light on whether there is a disparity (of the sort relevant to the Grand Illusion) between how our visual perception seems and the predicted nature of vision by informing us on how her vision seems to her. Thus, we need access to her pre-theoretical views to discover how vision really seems (without the influence of theoretical bias). This is, of course, a questionable assumption; it entails that we take her perception to be entirely void of theoretical bias, including bias from *folk theorizing*.¹⁸ As Dennett notes,¹⁹ there is strong evidence in support of

¹⁶ By theorizing (here), I mean the kind that the visual theorist is concerned with—she may be part of a community that engages in folk theorizing about perception.

¹⁷ *Surprise, Surprise*

¹⁸ I have benefited from discussions with John Dilworth on this matter. In my thesis, I discuss, at greater length, the possibility of an Ordinary Perceiver. I conclude that such an unbiased perceiver is not possible—at least *some* theoretical commitments are necessary in order to “know how to” maneuver oneself in order to achieve desired ends based on the information obtained through one’s perception (the fundamental thesis of the enactive approach).

¹⁹ Dennett, *How Could I Be Wrong? How Wrong Could I Be?*

the conclusion that the layperson *does* engage in folk theorizing. If she did not, then we would expect for her to not have *any* views about the nature of her own perception. But she does seem to have such views; what else could account for her confidence in her own phenomenology and her shock at the mere suggestion that it could be otherwise?

Nonetheless, for the sake of argument, let us grant that there exist Ordinary Perceivers of the variety that the debate favors—completely without commitments to the theorizing with which the Visual Theorist concerns herself²⁰—and that a perceiver without *these* commitments is sufficient for the proposed role of the Ordinary Perceiver in resolving the matter. Thus, we assume that the impact of the folk theory is insufficient to bias the experience of the Ordinary Perceiver in such a way as to destroy her function in determining whether she suffers from the Grand Illusion.

1.5 THE ORDINARY PERCEIVER'S BELIEF

We must figure out why she is surprised to discover her functional flaws—what belief was overthrown, what expectation unmet? Since the hypothesis is that it seems to the Ordinary Perceiver that her vision is everywhere photograph-like, her belief must also have the content that her vision seems “photograph-like” (and her behavior must be governed by expectations caused by this belief) in order for us to be warranted in claiming that she is subject to the Grand Illusion.²¹ This must be the belief that is overturned in order for her to have been subject to the Grand Illusion.

1.6 WHAT DOES (AND DOES NOT) SURPRISE THE ORDINARY PERCEIVER

²⁰ That is to say, at least without theoretical commitments to the functional flaws.

²¹ Due to the length of this paper, I omit my discussion of other potential explanations of why the Ordinary Perceiver is surprised by experiencing her functional flaws, such as other candidates for the belief (other than the belief argued for by Cohen).

Surprise is taken to be the indication by which we can determine how the Ordinary Perceiver's experience seems to her: it serves as empirical evidence that her experience does not accord with her expectations (which are generated by her beliefs). Since we must discover how her experience seems to her in order to determine whether she is subject to a Grand Illusion, it is crucial to address what does (and does not) surprise her.

This is particularly of interest to us, since it seems to be the prevalent method, within the debate, of attempting to establish that the Ordinary Perceiver is subject to the Grand Illusion. If we are unable to establish the Grand Illusion in this way, then those who argue in favor of it must find a new way of grounding their position. The interesting question then becomes one of how one could otherwise establish the disparity between how the Ordinary Perceiver's experience seems and how it actually is.²² Any new method of doing so must not allow any theoretical bias to enter into her interpretation of her experience—this rules out options such as asking the Ordinary Perceiver about her experience, because the language used in the questioning might bias her interpretation of how her experience seems to her.

As noted, Noë states that Dennett is one of the main proponents of the Grand Illusion and an advocate of the position that the Ordinary Perceiver's surprise is indicative of the disparity between the predicted nature of experience and how it seems to her.²³ Indeed, Dennett does appear to endorse such a view.²⁴ He argues that the surprise exhibited by the Ordinary Perceiver indicates that a belief has been overturned—that she was expecting something else. Since her surprise was caused by discovering her functional flaws, Dennett argues that she must have had a belief that

²² By which I mean the predicted nature of experience, which is assumed to be real nature of vision.

²³ Noë, *Is the Visual World a Grand Illusion?*

²⁴ Most notably in the following articles: *Seeing is Believing*; *Surprise, Surprise*; and *Filling In versus Finding Out*

led her to expect something else would occur: namely, the belief that her vision is everywhere detailed, distinct, continuous, in color, photograph-like, etc., which led her to expect that her vision would seem that way in this case, as it had seemed to her to be in all her past experiences.

In ‘Surprise, Surprise,’ a reply to Noë and O’Regan regarding the Grand Illusion, Dennett notes that his position disagrees with that of Noë and O’Regan with respect to the beliefs held by the Ordinary Perceiver. Noë and O’Regan argue that:

“But is it true that normal perceivers think of their visual fields this way [as in sharp detail and uniform focus from the center out to the periphery]? Do normal perceivers really make this error? We think not. . . . normal perceivers do not have ideological commitments concerning the resolution of the visual field. Rather, they take the world to be solid, dense, detailed and present and they take themselves to be embedded in and thus to have access to the world.”²⁵

Dennett argues to the contrary, stating that the Ordinary Perceiver *does* hold a belief that her visual experience is everywhere photograph-like. He states that such a belief can be exposed by observing the surprise exhibited by an Ordinary Perceiver during an experiment in which she discovers, say, her blind spot. He writes:

“Surprise is a wonderful dependent variable, and should be used more often in experiments; it is easy to measure and is a telling betrayal of the subject’s having expected something else. These expectations are, indeed, an overshooting of the proper expectations of a normally embedded perceiver-agent; people shouldn’t have these expectations, but they do. People are shocked, incredulous, dismayed; they often laugh and shriek when I demonstrate the effects to them for the first time. These behavioral responses are themselves data in good standing, and in

²⁵ Noë, *Is the Visual World a Grand Illusion?*

need of an explanation. They are also, of course, highly reliable signs of their “ideological commitments”—the very commitments that elsewhere in their essay the authors correctly cite as culprits that help explain resistance to their view... Surprise is only possible when it upsets belief.”²⁶

Dennett, as noted above, argues that the surprise exhibited by the Ordinary Perceiver indicates that her experience leads her to expect something other than the experience generated by the flaws. Thus, she is subject to the Grand Illusion.

Noë notes that, though the surprise discussed by Dennett must be explained, so must *lack of surprise*.²⁷ He notes that the Ordinary Perceiver is not surprised when she must put on her glasses, relocate, look around, squint, etc., to get better visual information. It seems that she would be surprised at the need for these adjustments if she believes that her vision is everywhere detailed, distinct, continuous, in color, photograph-like, etc. It is not clear why such adjustments would coincide with her expectations (thus preventing the surprise) if her experience seemed to her as though she had all the detail at once, like a photograph. He writes:

“Surprise requires explanation, but so does lack of surprise. Notice that we are not surprised or in any way taken aback by our need to move eyes and head to get better glimpses of what is around us. We peer, squint, lean forward, adjust lighting, put on glasses, and we do so automatically. The fact that we are not surprised by our lack of immediate possession of detailed information about the environment shows that we don’t take ourselves to have all of that information in consciousness at once. If we were committed to the snapshot conception, wouldn’t we be surprised by the need continuously to redirect our attention to

²⁶ Dennett, *Surprise, Surprise*

²⁷ Noë, *Is the Visual World a Grand Illusion?*

the environment to inform ourselves about what is there?"²⁸

This implies that the Ordinary Perceiver is not committed to the belief that she experiences "photograph-like" vision.

Taking stock, then, we can see that the Ordinary Perceiver *is* surprised to discover her flaws and is *not* surprised by her need to move around, put on her glasses, look around, etc., to obtain the desired visual information. This means that a satisfactory account must include an explanation for both the surprise and lack thereof. We must still determine the methods needed to develop a satisfactory explanation of this behavior.

2. THE PREDOMINATE METHOD WITHIN THE DEBATE FOR DETERMINING WHETHER THE ORDINARY PERCEIVER IS SUBJECT TO THE GRAND ILLUSION

2.1 OVERVIEW

Let us start at the beginning and work through the steps towards the conclusion that the Ordinary Perceiver is subject to the Grand Illusion (in accordance with the approach most commonly utilized within the debate). We notice that the Ordinary Perceiver is surprised to discover her flaw(s). We take her surprise to indicate that a belief has been overthrown, and we wonder what the content of that overthrown belief must have been in order for her experience of the flaw to have been surprising to her. Surprise only occurs when our experience fails to conform to our expectations. Our expectations are generated by our beliefs, which are, in turn, caused by our experiences. As such, surprise is taken to be our outward display of the overthrowing of the belief that generated the expectations which were unmet in the

²⁸ Noë, *Is the Visual World a Grand Illusion?* pg. 7

surprising situation. Her surprise, then, was caused by the failure of her experience in this case to conform to her expectations, resulting in the belief responsible for the faulty expectation being overthrown.

For example, suppose that she experiences her blind spot. In this case, her experience is such that a portion of her visual field fails to have incoming visual information. She exhibits surprise during this experience. However, from this point forward in our evaluation of the significance of the surprise, there are two ways of proceeding. At this point, I will follow the manner used in the Grand Illusion hypothesis, to see how it turns out. Later, I will return to this point²⁹ in an attempt to formulate a new interpretation of the significance of the surprise.

In order for the surprise to be generated by beliefs which she acquired through past experience (such that her belief that her vision seems "photograph-like" was generated by her experience being "photograph-like"), her past visual experiences must have generally seemed to be everywhere detailed, distinct, continuous, in color, etc. Those who argue for the Grand Illusion hypothesis must give an account of how it is that she fails to experience her flaws under normal conditions, such that she is able to have the requisite "photograph-like" experience which could generate the purported belief that her vision is "photograph-like."

However, her surprise alone is still insufficient to warrant attribution of the Grand Illusion to the Ordinary Perceiver. While it can be argued that the surprise exhibited by the Ordinary Perceiver when made aware of their functional flaws for the first time is indicative of a belief being overthrown, the debate needs a means of explaining and resolving the disagreement between those who argue that this is sufficient evidence that the Ordinary Perceiver is subject to the Grand Illusion and

²⁹ Chapter 4

those who argue to the contrary. The crucial distinction, I would suggest, lies in the method of belief acquisition: how did she come to have the belief in question? To this end, I propose a distinction between two ways that the relevant belief could have been acquired: experience-based beliefs versus inference-based beliefs.

2.2 DISTINCTION REGARDING METHODS OF BELIEF-ACQUISITION

Experience-based beliefs (call them ‘belief_e’) are acquired through perceptual experience. Inference-based beliefs are acquired through inferences made from other beliefs. In other words, I acquire an inference-based belief, (call these ‘belief_i’) as a result of inferring the content of that belief_i from another belief (which may itself either be a belief_i or a belief_e). An example will be helpful to illustrate the distinction. Consider the following incident which happened to me recently while shopping for a new shirt. I was studying a particular shirt, commenting to my companion that I liked the shirt very much except for the pattern. My companion then drew my attention to another shirt across the table, and asked whether I liked that shirt better. This shirt had a pattern that I liked much more. I decided that it was better than the other one, since, other than the more preferable pattern, it seemed to be qualitatively identical in all other respects. It was only after I had made that statement aloud regarding my judgment that the two shirts were otherwise alike that my companion to my attention to an interesting detail that I had missed: the shirt was not, as I had thought, similar in all other respects—the collar was entirely different (the first had a button-less v-neck collar, while the second was a polo shirt with three buttons).

My belief that this shirt had the same kind of collar as the first was inferentially acquired. Was I subject to illusion in this case? Though I judged the two shirts to have similar collars, it does not follow that I had illusory experience. I made my judgment about the second shirt without having any experience of the collar at

all—it was not the case that I had an *illusory* experience of the second shirt as having the same collar as the first (such that I would have acquired a belief_e that the collars were the same style); instead, that the collar seemed to be the same was a judgment that was not based on perceptual experience, illusory or otherwise.

The proposed distinction does important work for us regarding what is meant by ‘seems to be the case to the Ordinary Perceiver.’ Discussion of how something *seems* to an individual appears (at least in this debate) to refer to the *experience(s)* that are being had by the individual. The distinction discerns between cases where it *seems* to me (experientially) that the second shirt has the same kind of collar as the first, and cases where I infer that the second collar is the same as the first (perhaps I implicitly assumed that all of the shirts on that table must have the same kind of collar, and I failed to have any experiences of the shirt that suggested that it was not so). It is only in the experiential cases that we have the “seeming” that is required for the illusion. The beliefs_i, in contrast, only tell us that it is not the case that it seems to the Ordinary Perceiver that her vision is not everywhere detailed, distinct, continuous, in color, photograph-like, etc. This is so because she is not receiving the requisite *experiential* information to form a belief regarding a positive statement about how her experience seems to her; instead, the information that she *is* receiving (primarily from the foveal region) seems photograph-like and, given no information that it is otherwise, she *infers* that the rest of her visual field is this way as well.

It is important to note the distinction between what is meant by “seems” when it is used to describe the content of the two kinds of beliefs. Regarding the content of the experientially-acquired belief, by “seems”, we mean that, to the perceiver, the content of the perceptual experience was sufficiently robust to cause the belief. By “sufficiently robust,” I mean that the nature of the perceiver's experience is sufficient

to cause the belief (rather than relying on the kind of inference that led to my belief that the shirts has the same kind of collar). In my earlier example, the experientially-acquired belief that the second shirt had the same kind of collar as the first could only have been generated by having an (illusory) experience of the shirt as having that style of collar. In contrast, “seems”, as used to describe the content of an inferentially-acquired belief, means that the perceiver judged it to be so without having actually experienced it as such. For example, though I might have been surprised to learn that the second shirt had a different style of collar than the first, and I might have exclaimed the collars seemed to be the same to me, my belief that the second collar was the same style was not acquired through a perceptual experience of the second shirt as having a v-neck collar. Therefore, it seems incorrect to say that I was subject to an illusion—rather, I acquired the belief through the fact that the part of the shirt to which I was attending was similar in style to the first, and that I was not perceiving any information to the contrary; in fact, I had not glanced at the collar at all. The fact that I had not, at any point during which my belief about the similarity in style between the two shirts was formed, looked at the collar makes it very difficult for one to argue that I was subject to an illusion—I cannot be said to have experienced the collar of the shirt as a button-less v-neck (which would have been the illusory experience) if I did not experience the collar at all.

Thus, we can see the importance of the belief being acquired experientially. This difference between the content of these beliefs (though *functionally* equivalent)³⁰ matters because, in order to attribute to her the Grand Illusion, it must *seem* to the Ordinary Perceiver that her vision is everywhere photograph-like.³¹ This “seeming”

³⁰ By functionally equivalent, I mean that both beliefs would cause the Ordinary Perceiver to exhibit surprise when she discovers her flaws.

³¹ I suppose that it could be argued that inference-based beliefs could be sufficient to motivate the Grand Illusion hypothesis, but this seems to be a non-standard use of the term ‘illusion’. Illusions are understood to be experiential and, though the Ordinary Perceiver would exhibit the surprise if the

can only be established in the case of beliefs_e. Thus, we must determine whether this belief, if held by the Ordinary Perceiver, could have been experientially-acquired such that we are warranted in attributing the Grand Illusion to her.

2.3 THE METHOD OF ACQUIRING EXPERIENCE-BASED BELIEFS

How could she acquire the requisite belief experientially? Since she could only do so through experiences that seemed everywhere photograph-like, we must determine whether such an experience is possible. Note that there are two criteria that our account of her experience must satisfy: her experience must seem photograph-like during normal perception, but we must also account for how she is able to experience her flaws under non-normal conditions. One suggestion is that the functional flaws are “filled-in.”³² “Filling-in,” at least as it is understood with respect to this debate, refers to a theoretical process whereby the brain edits and adds information (hence “fills-in”) in the regions of the visual field that are missing visual information due to functional flaws—allegedly accounting for why we do not notice these gaps of missing information.³³ Suppose that we wish to entertain some variety of “filling-in.” This would allow for the Ordinary Perceiver to have the experience requisite for acquiring the belief. Would it also warrant us in attributing to her the Grand Illusion? It does not: her belief that her vision is everywhere detailed, distinct, continuous, in color, photograph-like, etc., would be *true*; her experience would not be illusory

experience failed to conform to her expectations as generated by an inference-based belief, it seems to be a misnomer to call her previous experience an illusion, in the same way that it seems incorrect to say that I suffered from an illusion in my “shirt collar” example.

³² I intend for this term to be understood, as used here, as referring to a theoretical process whereby the “gaps” that are a consequence of the flaws are touched up in such a way (the bit-map “filling in” of Dennett’s *Filling-in vs. Finding Out*) that the gaps are eliminated. We perceive the experience as “gap-free.”

³³ I contend that the argument that we ought to experience these “gaps of information” under normal circumstances (or, at least, that is what perception *really is like*) is invalid. This is because it does not follow from the fact that we can experience these flaws under certain circumstances that either (a) we ought to expect to experience them at all times, or (b) to not experience them under other circumstances is indicative of being subject to an illusion.

because the filling-in would make her vision seem that way because it *is* that way. In other words, according to filling-in, the perceptual product—the object of perception—is everywhere detailed, distinct, continuous, in color, etc., such that photograph-like perceptual experience is veridical. However, the challenge³⁴ for those who would argue for filling-in is to explain how it is possible for us to *ever* experience our flaws. Given that the nature of “filling-in” is such that that we do not experience the flaws because what we perceive is a filled-in perceptual product which is developed further down the causal chain (retinal image is touched up), it is not clear why filling in would fail us (since the image does not depend strictly on the incoming information—it is an edited version of the original information).

Perhaps an advocate of filling-in might object to my dismissal of her position as a viable explanation for the Ordinary Perceiver’s experience.³⁵ After all, doesn’t our failure to experience these flaws warrant *some* explanation—and what could be better than an explanation that argues that we would not expect to experience them because they are “filled-in”? We might inquire as to how the “filling-in” proceeds—that is, how it is that these flaws are edited. The advocate of a veridical variety of “filling-in” could answer, as so many have, that the brain “knows what goes there” and “fills it in” accordingly. However, this explanation has a bizarre, undesirable consequence: we would expect that, in unfamiliar environments, our functional flaws would initially be exposed until our brains figure out “what goes there.” Since we do *not* experience this occasional exposure of the flaws,³⁶ we can reject explanations that

³⁴ The need to meet such a challenge is removed once the Grand Illusion hypothesis is rejected in favor of my proposed reformulation of our expectations. See footnote 29, below.

³⁵ See my argument for how, once we have modified our expectations appropriately, the advocate of filling-in no longer owes any explanation regarding why filling-in “fails” when we perceive our flaws. Furthermore, all of the undesirable consequences (both here and elsewhere in the debate) are removed, such that filling-in can be seen as a viable theory of perception, rather than the oddity that it is taken to be at this point in the discussion.

³⁶ I can assert this with confidence, given that Ordinary Perceivers exhibit surprise upon experiencing their flaws, a reaction that we would not expect if they experienced them occasionally; it is the novelty

posit veridical filling-in³⁷

Suppose that one were to argue for a non-veridical variety of filling-in. The bizarre consequence of the veridical filling-in account would be avoided—but at what cost? One might argue that the content need not be veridical—any filling-in will do, just so long as the perceiver does not experience the gaps. However, this has a bizarre consequence of its own. Depending how this filling-in is selected, we might be *more* likely to notice our functional flaws: if the assignment of the filling-in design (what the filling-in looks like) is random, then we would expect our functional flaws to be obvious to us. The content of the filling-in could be different from the incoming visual information around it, such that we could more easily detect the flaws due to the anomalous information.³⁸ The advocate of non-veridical filling-in might counter by arguing that, since these flaws are not located in the foveal region, like the rest of the visual information that we acquire external to the fovea, we do not experience it in enough detail to notice the difference. However, this claim is not valid: while their assertion about not experiencing visual information acquired external to the foveal region as highly detailed is true, the same argument cannot be made for the contents of filling-in—unlike the visual information, which requires photoreceptors to be acquired, filling-in is argued to be generated by the brain, which is not subject to the same restrictions. One could, of course, argue that the brain represents that content as poorly as the visual information attained through that region external to the fovea (perhaps so as to not create disparities between the “filling-in” and the surrounding visual information), but this begins to look like an unnecessarily complex

of the experience which warrants the surprise (and which would be lacking if this explanation were true).

³⁷ At least insofar as the account of “filling-in” requires “knowing what goes there” (and, perhaps due to my unfortunate lack of imagination, I cannot fathom any other way of attaining veridical “filled-in” content).

³⁸ Rather than experience the flaws as poor or absent visual information, we would now experience them as containing anomalous information as compared to the rest of the visual field.

explanation.

Could our Ordinary Perceiver gain the experientially-based belief any other way? It seems that she could not. Her being able to acquire the belief through experiencing her vision as everywhere detailed, distinct, continuous, in color, photograph-like, etc., requires that her vision *is* this way (thus enabling her to experience it as such).³⁹ However, if her vision *is* this way, then she is not subject to any illusion—the content of her experience does not deviate from the predicted nature of experience⁴⁰ and her belief that her experience is “photograph-like” is true.⁴¹ Indeed, any experientially-based belief that is acquired through experiencing veridical, filled-in content will render (1) false; there is no discrepancy between the nature of our vision and how it seems to us. This rules out the experientially-based belief with the content ‘my vision is everywhere detailed, distinct, continuous, in color, etc.’ as a viable option.

2.4 OTHER CANDIDATES FOR THE BELIEF

In order to be justified in attributing the Grand Illusion to the Ordinary Perceiver, we must discover whether the belief that her vision is everywhere

³⁹ There is, perhaps, one other option, which I do not entertain here (because I do not see any reason to take it seriously): continuous hallucination. This would not require “filling-in” (this is different from the brain “touching-up” the contents, and could be entirely accidental) but would involve constant hallucination in the regions of those flaws.

⁴⁰ This is because the object of perception is, on such an account of filling-in, a touched-up image, which *is* everywhere detailed, distinct, continuous, in color, etc. Given that an account that uses filling-in as an explanatory resource must posit that perception is edited in this way, then our expectations would be altered from the original set in the debate—we would expect that our vision is everywhere detailed, distinct, continuous, in color, etc., because that which we perceive (the filled-in perceptual product) is that way.

⁴¹ Further, it is worth noting that it is unclear how we could *ever* experience the flaws (if they are “filled-in”). Perhaps the guise is only active when we are not attending to it or, like a poorly placed mask (think of enjoying a play from a seat where one’s angle to the stage is such that one only sees the sides of the masked actors—we would see what the person on stage *really* looked like (behind the mask), and not the mask—whereas the rest of the theatre enjoys the disguise), we are unable to see it under certain circumstances. Still, this begins to feel ad hoc—the more “just so” our account becomes, the more we may rightly long for a stronger explanation.

photograph-like, is the only belief that would cause her to exhibit such surprise and, if not, whether it is at least equally as strong of an explanation as the attribution other beliefs (regarding how well it fits the relevant information). As such, we must evaluate other candidates for the belief that would cause her to exhibit this surprise. If any of these candidates are viable, we must determine whether the belief_e that her vision seems photograph-like has at least comparable, if not greater, explanatory strength.

To illustrate that viable alternatives exist,⁴² we will address one such alternative raised by Jonathan Cohen.⁴³ He suggests that, rather than having theoretical commitments regarding the nature of her visual field (which he believes that we are not warranted in attributing to her), the Ordinary Perceiver simply believes that they notice objects and events that are located or occurring in front of them. Such a belief would also provoke surprise in light of experiencing the functional flaws—it would generate expectations for the Ordinary Perceiver that (it seems to her as though) she never fails to notice objects and events that are directly in front of her. The Ordinary Perceiver would be surprised upon discovering her blind-spot (where part of the object(s) are occluded), experience change-blindness or inattentional-blindness (failing to notice events happening before them due to lack of attention) because these experiences fail to conform to the expectations that would be caused by such a belief.

As such, the requisite belief is not the only one which would elicit surprise in

⁴² Due to the length of this paper, I omit my discussion of other potential explanations of why the Ordinary Perceiver is surprised by experiencing her functional flaws, such as other candidates for the belief. The importance of illustrating that there is at least one other candidate for the belief serves as evidence that we must continue our inquiry; that the Ordinary Perceiver is surprised—and that the surprise *can* be explained by the requisite belief—is insufficient grounds for attributing to her the Grand Illusion. However, see also my later discussion of beliefs about the nature of occlusion as another potential explanation.

⁴³ Cohen, *The Grand Grand Illusion Illusion*

the Ordinary Perceiver. Since this is so, we must further examine the beliefs that we could attribute to the Ordinary Perceiver to determine which is the best explanation of her surprise (and whether the surprise is indicative of the Grand Illusion).

2.5 FURTHER CONSIDERATION OF INFERENCE-BASED BELIEFS

If I am correct in concluding thus,⁴⁴ then the only candidate (with the content ‘my vision is everywhere detailed, distinct, continuous, in color, etc.’) for the cause of the Ordinary Perceiver’s surprise is a belief_i. This belief_i would cause her to exhibit surprise upon discovering the functional flaws of her visual system. The Ordinary Perceiver would have the tacit belief⁴⁵ that it is not the case that her vision seems to her as though it is not everywhere detailed, distinct, continuous, in color, photograph-like, etc. One might wonder why I assert that beliefs_i are held *tacitly*. Arguably, not all our beliefs are held overtly; some of our beliefs, though we may not actively entertain them or have the ability to express them linguistically (or even know that we have these beliefs), nonetheless impact our behavior. For example, though an Ordinary Speaker of the English language may not have explicit access to her beliefs about the conjugations of words in the English language (and, given what I know about the teaching of the English language in most schools, it is safe to assume that most Ordinary Speakers of English do not), she can nonetheless detect the occurrence of a conjugation error. This is similar to what Pylyshyn refers to as ‘tacit knowledge’: though the individual might not be able to make use of such information for the purposes of answering questions about the matter, she can reveal such knowledge by

⁴⁴ That there is no viable account that would explain how she could experientially-acquire the requisite belief while nonetheless still be able to experience her flaws under non-normal perceptual circumstances.

⁴⁵ By ‘tacit belief,’ I mean roughly what Pylyshyn describes as ‘tacit knowledge,’ though I do not intend to imply any further similarities between the ‘mental imagery debate’ and the Grand Illusion at this point.

exhibiting telling behavior when occurrences are contrary to her knowledge.⁴⁶ Likewise, this would account for both the surprise exhibited by the Ordinary Perceiver upon experiencing the flaws *and* her strange looks and responses when queried about whether she believes that their vision is everywhere detailed, distinct, continuous, in color, photograph-like, etc.⁴⁷

Perhaps one might object that my explanation is incomplete yet; I have not yet shown how inference-based beliefs do not warrant our attribution of the Grand Illusion to the Ordinary Perceiver. After all, it appears that the content of the inference-based belief is at odds with the “real” nature of perception.

To explain how inference-based beliefs do not warrant the attribution of the Grand Illusion to the Ordinary Perceiver, it is useful explicating two seemingly disparate accounts by Noë and Dennett, illustrating how these accounts can be interpreted as sharing much common ground. This common ground illustrates that the surprise, if generated by inference-based beliefs, is insufficient to warrant attribution of the Grand Illusion to the Ordinary Perceiver.

In an attempt to explain why the Ordinary Perceiver is surprised by her functional flaws, Noë introduces the problem of *perceptual presence*.⁴⁸ This problem is a strengthened version of the Grand Illusion hypothesis. The worry is that we seem to “experience” a whole object when we are, in fact, only perceiving part of it. For example, when we close our eyes and hold a bottle in our hand, it seems to us as though we are experiencing the whole bottle when, in fact, our hand is only making contact with a portion of it. Likewise, when we see a cat on the other side of a fence, though it seems to us that we experience the whole cat, we only perceive part of it

⁴⁶ Pylyshyn, *Mental Imagery: In Search of a Theory* (particularly section 3.1)

⁴⁷ Ordinary Perceivers generally reply to inquiries about their visual experience with either puzzled looks or thoughtful replies stating that either they do not believe their experience to be that way or that they had never thought about it.

⁴⁸ Noë, *Is the World a Grand Illusion?*

(that which is visible between the boards of the fence).

Noë argues that the key to resolving the problem of perceptual presence is through analysis of amodal perception. Amodal perception is, “perceiving what is, strictly speaking, out of view.”⁴⁹ Just as we amodally perceive the hidden parts of the cat, Noë suggests that we amodally perceive objects in our environment which we are not currently *visually* experiencing. He states that this proposed alteration to the problem changes the question from one of whether the Ordinary Perceiver has the (illusory) experience as though she perceives everything in the environment at once to one of whether amodal experience is illusory. Noë argues that there is no need for the brain to create a representation of, say, the whole cat (which is partially occluded by the fence) since all we need to do to get more information about the cat is to move around so that different parts of the cat become visible. We are skilled perceivers; we know how to manipulate objects and adjust ourselves such that we are able to perceive the information that we seek. When we squint or peer, as Noë puts it, we do so because that which we wish to see is yet too far or too small from our current visual perspective. We know what our range of visual clarity is and we adjust accordingly. Too big? Back up. Too small? Lean in.

At the end of the day, though Noë and Dennett have argued against one another’s positions regarding whether the Ordinary Perceiver is subject to the Grand Illusion, it appears that they have much in common. Both argue that the Ordinary Perceiver does not *experience*⁵⁰ vision that seems everywhere detailed, distinct, continuous, in color, photograph-like, etc.; rather, their experience is a matter of

⁴⁹ Noë, *Is the World a Grand Illusion?*

⁵⁰ Though Dennett claims that it does *seem* that way to the Ordinary Perceiver—contrast with Noë, who argues that it does not seem that way to her (the phenomenology involved in the Grand Illusion hypothesis is in error with respect to how it seems to her). Noë argues that, rather than perceiving a line which extends through the region of your blind spot (while you are attending to that region) as being unbroken, it seems to her as through her experience of the line is that it does not seem to be broken

representation, interpretation, or skilled (amodal) perception of those regions that fall outside of the region to which the perceiver is currently visually attending. Both accounts rely on a judgment sense of “seeming” (amodal perception, “finding out”), not on the experiential sense of the term. Understood in this way, both of their positions are consistent with arguing that the relevant belief held by the Ordinary Perceiver is an inferentially-based⁵¹ belief. The ‘amodal perception explanation’ offered by Noë appears to be very similar to the account that Dennett gives as an explanation of our experience of Andy Warhol’s Marilyn Monroe wallpaper—both argue that we do not *experientially* represent the poor or absent information as being present. Dennett writes:

“Suppose you walk into a room and notice that the wallpaper is a regular array of hundreds of identical sailboats, or—let’s pay homage to Andy—portraits of Marilyn Monroe. We know that you don’t foveate, and don’t have to foveate, each of the identical images in order to see the wallpaper as hundreds of identical images of Marilyn Monroe. Your foveal vision identifies one or a few of these and somehow your visual system just generalizes—arrives at the conclusion that the rest is “more of the same.” We know that the images of Marilyn that never get examined by foveal vision cannot be identified by parafoveal vision—it simply lacks the resolution to distinguish Marilyn from various Marilyn-shaped blobs. Nevertheless, what you see is not wallpaper of Marilyn in the middle surrounded by various indistinct Marilyn-shaped blobs; what you see is wallpaper composed of hundreds of identical Marilyns. Now it is

⁵¹ Presumably, Noë would not use the term “inference.” However, I only intend to imply that both his account and Dennett’s “finding out” involve something like a judgment. For example, on both accounts the belief that vision seems “photograph-like” is a judgment about the nature of vision based on the parts of the visual field that the perceiver can experience (and, from that, she judges that the flaws contain similar visual information), and not because she *experiences* the flaws as being “filled-in”, in the sense that there is some perceptual product which she *experiences* as “photograph-like”.

a virtual certainty that nowhere in the brain is there a representation of the wall that has high-resolution bit-maps that reproduce, xerox-wise, the high-resolution image of Marilyn that you have foveated. The brain certainly would not go to the trouble of doing that filling in! Having identified a single Marilyn, and having received no information to the effect that the other blobs are not Marylins, it jumps to the conclusion that the rest are Marylins, and labels the whole region "more Marylins" without any further rendering of Marilyn at all. Of course it does not seem that way to you. It seems to you as if you are actually seeing hundreds of identical Marylins. And in one sense you are: there are, as I said, hundreds of identical Marylins out there on the wall, and you're seeing them. What is not the case, however, is that there are hundreds of identical Marylins represented in your brain. Your brain just represents that there are hundreds of identical Marylins."⁵²

Notice another similarity between their accounts: they both offer interpretations that fit the behavior of the Ordinary Perceiver. Both accounts are consistent with the surprise (due to experiences that are contrary to her expectations) and lack thereof (due to experiences that conform to her expectations) that is exhibited by the Ordinary Perceiver. Neither account attributes to her expectations (or beliefs) that would lead her to fail to be surprised when experiencing the flaws—her surprise is consistent with both accounts, and the same is true for her lack of surprise. Thus, the assertions of each regarding what does and does not surprise the Ordinary Perceiver fit equally as well with the other's account as their own.⁵³

I will now consider whether the content of the inference-based belief is such

⁵² Dennett, *Filling In versus Finding Out*

⁵³ This is particularly interesting when one considers that the role that Dennett's assertions regarding what *does* surprise the Ordinary Perceiver and Noë's assertions about what does *not* surprise the Ordinary Perceiver each appear to be intended as evidence against the other's account. However, in this light, it appears that both sets of assertions fit *each* account equally well.

that we are warranted in attributing the Grand Illusion to the Ordinary Perceiver. Recall that the inference-based belief in question is held tacitly. If the Ordinary Perceiver tacitly believes that it is not the case that her vision seems to her as though it is not everywhere photograph-like then her belief does not entail that she also believes that her vision *is* everywhere detailed, distinct, continuous, in color, photograph-like, etc. No beliefs that assert positive claims are made about how her vision seems to her—only a negative one: it isn't the case that her vision is not photograph-like. She could acquire this negative belief because she has no reason to accept its negation: that her vision is not everywhere photograph-like. In other words, in order for her to acquire the positive belief that her vision is not everywhere photograph-like, she must have an experience that causes her to have that belief (such as experiencing the flaws). In the absence of such experiences (and experiences that would cause the belief that her vision is everywhere photograph-like—such as those discussed above), she is committed to neither belief and instead acquires the tacit belief that it is not the case that her vision seems to her as though it is not everywhere photograph-like. This tacit belief roughly amounts to the expectation discussed by Cohen: the Ordinary Perceiver believes that she perceives the events and objects that are directly before her.⁵⁴ In other words, her expectation is governed by her tacitly held belief that it is not the case that her experience is not everywhere photograph-like (she does not expect that she will fail to experience some of the objects and events immediately before her).

2.6 WHY DENNETT CANNOT ARGUE THAT THE ORDINARY PERCEIVER IS SUBJECT TO THE GRAND ILLUSION

Recall my earlier argument that the belief which is allegedly overthrown when

⁵⁴ Cohen, *The Grand Grand Illusion Illusion*

the Ordinary Perceiver is surprised at the experience of her flaws must satisfy two criteria: (a) it must have the content that vision seems to be everywhere “photograph-like,” and (b) this belief must be acquired through experience (her vision must seem “photograph-like,” in the experiential sense of seeming). We have seen that one cannot argue for the Grand Illusion by using the surprise as evidence unless the belief in question was acquired through experience. Now we will look at why one cannot make the case for the Grand Illusion unless one argues that beliefs have determinate content.

As noted earlier, Noë argues that Dennett is the main proponent of the Grand Illusion Hypothesis. However, I contend that he cannot argue for the Grand Illusion in the way that he attempts to do so. Since he is taken to be a central figure in the debate, it is worth reviewing the problem with his account which prevent him from being able to argue for the Grand Illusion, as well as one regarding his explanation for why the Ordinary Perceiver fails to experience her flaws.

There are two problems with Dennett’s account of the Grand Illusion hypothesis: the first is regarding how he attempts to establish his case for the Grand Illusion, and the second is regarding whether his “labeling” violates his Thrifty Producer Principle. Since Dennett’s account of how he attempt to establish his case for the Grand Illusion hypothesis is similar to other accounts that we have discussed which support their claim by noting the surprise exhibited by the Ordinary Perceiver, we are already familiar with that strategy. However, Dennett’s case is special, in that his holism appears to prevent him from being able to establish the Grand Illusion hypothesis. Once we have seen why this is so, we may then inquire as to whether one aspect of his account regarding why we do not experience our flaws violates a central principle in his general position.

Dennett argues that the surprise exhibited by the Ordinary Perceiver when she first experiences her flaws is indicative of a belief being overthrown—a belief that her vision is everywhere detailed, distinct, continuous, in color, etc. However, Dennett, unlike other advocates of the Grand Illusion hypothesis, does not mean that she has a belief with such determinate content; as a holist, he argues that beliefs have indeterminate content. It seems consistent with his overall position to interpret him as attributing this belief to her as part of an intentional stance that he adopts towards the Ordinary Perceiver. As such, the belief is attributed as a means of explaining her behavior.

This seems right so far as it goes—if she had such a belief, *then* we would expect her to exhibit surprise when she first experiences her flaws. The interesting question is whether there is reason to think this is the *best* stance to adopt towards her behavior. Important considerations regarding whether this is the best explanation⁵⁵ include how well it integrates with the rest of the beliefs that we attribute to her when we have adopted stance with the strongest explanatory power. In other words, are there beliefs that we attribute to her which would have stronger explanatory power regarding her behavior, as well as fit better with the rest of the best stance that we adopt towards her?

Compare this belief to one regarding the nature of occlusion. For example, it seems that we could explain her surprise does not indicate when she discovers her blind spot by attributing to her a belief about the nature of occlusion (as supported by her experiences thus far)—namely, the belief that an object (or some part thereof) is occluded from view if and only if it is blocked by another object. Surely, to one who

⁵⁵ Since his Intentional Stance is a predictive and explanatory strategy, meaning that we adopt a particular stance both to make sense of previous behavior and as a means of predicting future behavior (by hypothesizing how an ideally rational system might behave).

holds such a belief, it would seem surprising for an object to be occluded from view without being hidden behind another object. However, such surprise does not indicate that she is subject to the Grand Illusion; though her experience in this case fails to conform to her expectations, it is this, and not an illusory experience, that generates the surprise. As such, it appears that there are other beliefs that we can attribute to her which explain the surprise but which fail to indicate that she is subject to illusion.

Regarding which belief we should attribute to her, we need to consider which best fits with the rest of the intentional stance that we adopt towards her. While the former belief fails to connect up with any other beliefs we can attribute to her (perhaps this is why we are so resistant to saying that the Ordinary Perceiver really does have such a belief—it is so unlike the other beliefs that we could attribute to her that we have a hard time attributing such a belief), it appears that the latter could easily fit within the framework of cognitive development. For instance, if we assume that she acquired her belief about the nature of occlusion through previous experience, namely by having noticed that the only time that objects disappear is when they are hidden behind other objects, then she would not expect for an object to be occluded in that region of visual field (given that there is no object there which would occlude it).

Regardless of whether we can determine which belief best fits with intentional stance that has the greatest explanatory and predictive power, if attributing a belief with this content to the Ordinary Perceiver is necessary for making the case that she is subject to the Grand Illusion, then we ultimately end up with a skeptical problem. To understand how this problem arises we must first get clearer on how, exactly, Dennett attempts to make his case. Once we see this, it will become apparent that, in light of his holism, Dennett's position entails that there is no fact of the matter whether the

Ordinary Perceiver is subject to the Grand Illusion.

Recall that the importance of belief is merely to reveal the content of the Ordinary Perceiver's visual experience. When she exhibits surprise, we can infer that her experience failed to conform to the expectations generated by her beliefs. Thus, the surprise she exhibits upon first discovering her flaws is taken to indicate that a belief with the contents which would generate expectations which are contrary to her experience of flaws, which was generated by past experience, is overthrown.

Holism, however, entails that no beliefs have determinate content; the content of beliefs are indeterminate, meaning that there is no fact of the matter what is the content of belief. This is problematic because, as we saw earlier, the belief with requisite contents is necessary for establishing the Grand Illusion hypothesis in this way. Regardless of which belief we attribute to her, Dennett must⁵⁶ conclude that there is no determinate content to her belief, such that it is indeterminate whether she is subject to such illusion.⁵⁷

Regarding the second problem with his account, Dennett's explanation for why we fail to experience our functional flaws has two parts: (1) the brain ignores the missing data and (2) the brain "finds out" what information falls within the areas of missing data, and labels that region accordingly. Central to his argument that we ought to reject filling-in in favor of his account is his Thrifty Producer Principle—"if there's no one to look at it, don't bother making it". In other words, if there is no part of the brain designated to processing such information, it is superfluous for such information to be generated (presumably by whatever mechanism would be

⁵⁶ In light of his holism.

⁵⁷ Note that I am not arguing about whether beliefs have determinate content. I am only arguing that, given the use of the surprise as the means of determining whether the Ordinary Perceiver is subject to the Grand Illusion, the belief in question must have determinate content to make the case in this way. But this, of course, is independent of my own position regarding whether beliefs are determinate, which is irrelevant here.

responsible for the filled-in information). However, this principle in conjunction with the two parts of his explanation for why we fail to notice her flaws makes finding out superfluous as well as a violation of principle.

His account begins with arguing the brain simply ignores the missing data. He maintains that there is no part of the brain dedicated to processing visual information from the regions the visual field which, due to the perceptual flaws, lack such information to some extent. Since there is no part of the brain assigned to the task, the missing information goes unnoticed.

It is this claim that there is no part of brain assigned to the task of processing what would be the visual information from the regions which contain the flaws which motivates his Thrifty Producer Principle; that there is no part of the brain assigned to task seems (to Dennett, anyway) to entail that there is no part of the brain which would perceive the touched-up perceptual product, making any such product superfluous. Therefore, he concludes that there must not be any filling-in.

He then goes on to argue that the brain “find out” what falls under those regions of the visual field which are missing information due to the flaws and labels that region accordingly. However, this part of the account is problematic for two reasons. The first is that it is superfluous, given that he has already argued that the brain ignores the missing data. Secondly, it violates the thrifty producer principle.

An example of how “finding out” works will be useful here. Recall the earlier quoted passage in which he illustrates this point through asking the reader to imagine that she is standing before a wall full of Marilyn Monroe image wall paper by Andy Warhol. He argues that, though it seems to the Ordinary Perceiver as though she is seeing hundreds of Marilyn images, is not the case that, somewhere in her brain, there is a representation of hundreds of Marilyn images; rather, from the images that she

does see before her, and no disconfirming information to the contrary from the periphery, the brain just “jumps to the conclusion” that it is more of the same, and labels that region “more Marilyn images”.

Finding out is unnecessary given that he has already argued that the brain ignores the missing data. If the missing information is ignored, then it is not clear why we need to posit the labeling process. More specifically, it is not clear what work the labeling is doing. The labeling appears to be intended to play an explanatory role in his account regarding why we do not notice the flaws, but the ignoring feature of his account appears to fully explain this independently of the labeling. Furthermore, it is not clear how we could ever experience the labeling, since this seems like a perceptual product, and Dennett has argued that there is no part of the brain dedicated to perceiving such products.

Also problematic is that labeling appears to be at odds with the principle. As we have just seen, he argues that there is no perceptual product. However, labeling does appear to be just such a perceptual product. It is unclear what work labeling could do if we fail to have the mechanisms necessary to experience it.

Give these considerations, it seems that the best way to preserve the most of his account is to simply reject the labeling feature. When we exclude the labeling feature, the account is no longer internally incoherent. The ignoring feature and the principle are consistent with one another. Furthermore, this allows him to maintain his Thrifty Producer Principle, which plays an important role in his attack on the Cartesian Theatre conception of mind.⁵⁸

⁵⁸ This is one of the central moves of his position in general—he has been a long-standing opponent of such a theory of mind. (See ‘Consciousness Explained’, for example)

3. REEVALUATING THE ASSUMPTIONS THAT MOTIVATE THE GRAND ILLUSION HYPOTHESIS

3.1 WHAT ASSUMPTIONS MOTIVATE THE HYPOTHESIS?

Since our attempt to give an account regarding how it is that the Ordinary Perceiver's vision seems photograph-like under normal circumstances, but allows her to experience her flaws under special circumstances, seems fruitless, it will be useful to critically review the steps that lead us to this point. It seems that if we could somehow show that our conception of predicted (non-illusory) nature of perception was such that, if our expectations were different (that is, such that we wouldn't expect to experience the flaws under normal conditions), then we would no longer have the burden of explaining why we don't experience them normally but do experience them under certain circumstances. In other words, since the Grand Illusion hypothesis is motivated by the expectations generated by our assumptions (regarding the affect that the flaws would have on our perceptual experience), by reevaluating these assumptions we may find a way of altering our expectations, such that we no longer need to offer an account as to how this illusory experience of failing to experience the flaws is possible.

Recall that Dennett assumes that the surprise exhibited by the Ordinary Perceiver when she discovers her flaws indicates that she was expecting that her experience would be different from that what she actually experienced in this case. From this, he appears to assume that, since she was surprised by her experience of the flaws (the experience of which fails to be photograph-like), then she was expecting her experience to seem photograph-like, and she had such expectations because her previous experience had seemed photograph-like to her. However, this conclusion does not follow. It seems that it is reasonable to infer that her surprise only indicates

that she was not expecting to experience what she had when she discovered a flaw. In other words, since she was surprised to experience her vision as failing to seem photograph-like, all we can infer is that she was not expecting for her experience to not seem photograph-like, but from that it does not follow that she did expect to experience to seem photograph-like.

One might wonder why I argue that Dennett implicitly assumes (or, at the very least, *inadvertently* assumes, such that it is entailed by his position) that she is expecting that her experience will be photograph-like. The following seems to be a reconstruction of how he determines her expectation. Consider the role that the surprise performs in determining whether the Ordinary Perceiver is subject to the Grand Illusion. Since she is surprised to discover her flaws, Dennett infers that she was expecting something else, namely that her experience would be photograph-like. This expectation must have been caused by a belief that her experience would be so—and this belief must have been generated by previous experience(s) which did seem photograph-like to her (at the time that the belief was formed and throughout the duration which she held the belief). He happily notes that her experience did seem photograph-like to her, and he concludes this by taking as evidence the surprise that she exhibits when she experiences her flaws for the first time. The only way that his conclusion (that her previous experience seemed photograph-like to her) follows from the surprise is with the inclusion of the missing premise that the overthrown belief is that her vision seems photograph-like, such that when she experiences her vision as not-photograph-like (during the discovery of her flaws), it is *that* belief that is overthrown. Otherwise, all he would be warranted in concluding from the surprise is that she was not expecting the not-photograph-like experience, but this belief could be generated by more than just the photograph-like experience. Whereas, the only belief

(overthrown upon discovery of the flaws) that would support his claim that her experience seemed photograph-like would be the experientially-acquired belief that her vision seems photograph-like, such that the “something else” that she was expecting was that her experience of the region of, say, her blind spot, would not be not-photograph-like (since she experiences that area during normal perception as photograph-like).

The surprise exhibited by the Ordinary Perceiver could be interpreted as indicative of both (by which I mean her expectation that her experience would be photograph-like or that it was not the case that she expected her experience to not be photograph-like would both cause surprise). However, in the case that she was not expecting that her experience would not be photograph-like, the Ordinary Perceiver cannot be said to suffer from the Grand Illusion, because we can only establish that it is not the case that her vision seems to her as though it is not everywhere photograph-like. But this is different from seeming to the Ordinary Perceiver that her vision *is* everywhere photograph-like. These beliefs are functionally equivalent (they cause the Ordinary Perceiver to behave in the same way—at least regarding the display of surprise), but their content is distinct. However, from the fact that it does not seem to the Ordinary Perceiver that her vision is not everywhere detailed, distinct, continuous, in color, etc., it does not follow that her vision does seem to her to be everywhere detailed, distinct, continuous, in color, etc. Her vision could appear many different ways and still satisfy the criterion ‘does not seem to her as though her vision is not everywhere detailed, distinct, continuous, in color, etc.’ Only if we can make the case that she has the belief_c that her vision seems everywhere detailed, distinct, continuous, in color, etc., are we warranted in inferring that it does seem to the perceiver as though her vision is everywhere photograph-like.

Consider this analogous case: our perceiver is walking through her local campus, when she gets hit with a water balloon. She is surprised by this, but why? Is it because she is expecting to not get hit with a water balloon? Presumably not--she could be surprised by this event even if she has no previous experience with water balloons (such that she is not aware that such things even exist). If surprise was only possible in cases where one was expecting the exact opposite event from the one experienced (in this case, if the perceiver could only be surprised if she was expecting to not get hit with a water balloon), then the only circumstances under which one could be surprised would be those which one is both aware of the source of the surprise (minimally, that they know such a thing exists) and expecting that their experience will be such that the source of the surprise will be absent. However, this entails that, by definition, the Ordinary Perceiver is not surprised by her experience of the flaws. This is because, due to the criterion of being unacquainted with the flaws as a prerequisite for being Ordinary Perceiver, she does not know that the source of her surprise exists, such that she could not possibly expect not to experience it.

One might protest that the advocates of the Grand Illusion hypothesis never argue that she was expecting the exact opposite of her experience and, as such, my claim that it is entailed by their position is unjustified. While it is true that no party explicitly argues as such, my conclusion is still entailed by their position. One can see this by looking at how they interpret the surprise. From the Ordinary Perceiver's surprise, they argue that she was expecting something else—which ultimately turns out to be a continuation of her previous, "photograph-like" experience. What is of interest here is how they get from her surprise to the conclusion about the "photograph-like" quality of her experience. Since she was surprised because her experience did not conform to her expectations (which are generated by a belief that is

overthrown when she experiences her flaws for the first time), it appears to be implicitly assumed that her belief was caused by her experience being "photograph-like" due to the widespread discussion regarding why it is that she fails to experience her flaws. To put the point another way, it is unclear why there would be so much debate over how we should account for her failure to experience the flaws if it is not assumed that, under normal conditions, she perceives the area which contains the flaws (but nonetheless fails to experience them). And, of course, I am not arguing that the Ordinary Perceiver analyzes her experience in this way. But this is in contrast to the belief that is requisite for warranting our attribution of the Grand Illusion to the Ordinary Perceiver. If it is not the case that her vision seems to her as though it is not everywhere detailed, distinct, continuous, in color, photograph-like, etc., this is different from seeming to her that her vision *is* everywhere detailed, distinct, continuous, in color, photograph-like, etc. The inference-based belief does not include any positive assertion about what her experience *is* like (only what it is *not* like); numerous different positive assertions about the nature of her experience are consistent with this tacit, inference-based belief. Since this belief does not have the requisite content, we cannot attribute the Grand Illusion to her. The onus is on the advocate of the Grand Illusion to give an account of how the Ordinary Perceiver can experientially acquire the belief that her vision is photograph-like that does not run into the problems already raised.

Given that these expectations are what motivate the claim that the Ordinary Perceiver is subject to an illusion, it seems plausible that by changing our assumptions, we may be able to eradicate the problem. To figure out how to change our expectations, we must carefully determine what they are, and see what assumptions led to these expectations. Once we have done this, we can challenge

those assumptions to determine whether we ought to adopt new assumptions which would alter our expectations in the desired way.

3.2 A DISTINCTION BETWEEN KINDS OF FUNCTIONAL FLAWS

It seems that the best place to start our evaluation of the assumptions that motivate the Grand Illusion hypothesis is right back at the beginning, with the functional flaws. As noted earlier, there are many different flaws, the failure to experience all of which is taken to be indicative of the Grand Illusion. However, we can make a useful distinction between two kinds of flaws based on the features of the visual system which generate them.

As noted above, part of our explanation must include attention to the predicted nature of perception. I believe that the predicted nature of perception, as understood in the Grand Illusion, rests on an invalid argument. Recall that, by the predicted nature of experience (as formulated within the Grand Illusion hypothesis), I mean the nature of perception as including the experience of the functional flaws. Given that it is assumed that perception that fails to include experience of the flaws is illusory, then presumably, if our experience included the flaws, we would not be subject to an illusion. This is derived from an argument like the following:

4. We have functional flaws, which we can experience under certain conditions.
5. We do not experience these flaws under normal conditions.
6. Since the Ordinary Perceiver has not yet experienced these flaws (and, upon doing so for the first time, exhibits surprise), we may infer that she was expecting something else to happen—thus overthrowing her belief that caused her expectation to the contrary (that her vision is detailed, distinct, continuous, in color, photograph-like, etc.). She is subject to an illusion—she does not experience her

vision as it *really is*.

However, it does not follow from the fact that we have these functional flaws which we are able to experience under certain conditions that we must notice them under other or all conditions. Furthermore, it does not follow that the predicted nature of experience which motivates our expectation that we would experience our flaws under normal conditions is based on the best possible set of assumptions that could be formulated based on the data from the relevant perceptual research.

This appears to be a problem with our expectations: because we *know* that the flaws exist, we *expect* that we should experience them. However, it is not clear that we are warranted in having such expectations. As noted above, from the fact that we have these flaws which we can experience under certain conditions, it does not follow that we ought to experience them in any other or all circumstances.

As we saw, the Grand Illusion hypothesis entails an incoherent formulation of the predicted nature of perception. Though never explicitly stated, this formulation can be directly derived from that which is taken to be illusory about the Ordinary Perceiver's visual experience. The error ultimately stems from considering the data from the two different kinds of flaws independently, then formulating expectations for each kind without considering the data of the other. When we simultaneously evaluate the two contrary accounts of predicted nature of perception, it is apparent that the formulation is incoherent.

A distinction can be made between two kinds of functional flaws. The differentiating feature between the two kinds is the aspect of the perceptual system that causes them. One kind is generated by the functional structure of the visual system. Let us call these the 'structural flaws'. The other kind is generated by the attention-dependent nature of perception. We will call these 'scope flaws', as these

flaws restrict the scope of that which we notice about the changing world before our eyes.

The structural flaws are caused by the functional structure of our visual system. By functional structure, I mean how the structure of our visual system impacts the quality of our visual information. Perhaps the best example of this kind of flaw is blind spot. The blind spot is created by the lack of photoreceptors in front of where the optic nerve connects to the back of the eye. Since there are no photoreceptors, we fail to receive visual information from the region.

Other examples of the structural flaws include the saccades, uneven distribution of rod and cones (rods primarily external to foveal region, cones restricted to foveal region) and a lessening in general of photoreceptors from the center of the visual field (where there is a high density of cones) towards the periphery (where there are no cones and very few rods).

It is important to understand the impact of the structural flaws both on her visual experience under normal conditions and also under contrived conditions which do not conform with our normal perception. The discrepancy between how the flaws impact our vision under normal circumstances and under special conditions leads the advocates of the Grand Illusion hypothesis to argue that our experience under normal circumstances is illusory. The assumption appears to be that, since we know that we can experience these flaws, we must be subject to illusion when we fail to experience them. In other words, since the functional structure of our visual system is responsible for her visual experience, then one must be subject to illusion if one's visual experience fails to conform to that which the functional structure would support. Thus, it is useful to discuss the effect each flaw has on our visual experience under both kinds of circumstances, according to the Grand Illusion hypothesis.

As noted earlier, the blind spot results from a lack of photoreceptors. Under normal conditions, we fail to notice our blind spot. However, we can easily experience our blind spot under contrived conditions, such as shutting one eye while fixing the gaze of the other on some fixed point, then attending to that region of the visual field (this usually requires some sort of image or symbol which we can observe to be altered by the blind spot—disappearing behind blind spot, etc.).

The uneven distribution of rods and cones results in differing quality of visual information from different areas of the retina. Cones are the photoreceptors that allow for color vision. They only work well in well-lit environments; they are functionally impaired in poor lighting. The high density of cones in the foveal region creates a highly detailed, richly colored visual experience of the center of our visual field. However, we have few rods in the area, making it very difficult to acquire visual information from that region in low lighting conditions.⁵⁹ Outside of the foveal region, we fail to have color vision. This is because rods, which are responsible for our ability to see under poor lighting conditions, do not support color vision.

Given the uneven distribution of, and differences in kind of visual information processed between, the two kinds of photoreceptors, the advocates of the Grand Illusion hypothesis conclude that our visual experience under normal conditions, which we do not experience as lacking color and detail in areas, must be illusory. Note that we are able to experience this kind of structural flaw under contrived circumstances. For example, we cannot detect the color, and sometimes even the shape, of an object that is held at our periphery. At best, we can detect it if it is moving—otherwise, we seem to be unable to discriminate anything about it.

Unlike the functional flaws, the scope flaws are not generated by the physical

⁵⁹ This is why, for instance, you cannot see most stars by looking directly at them. Instead, you must look slightly to the side of them.

structure of the visual system; they result from the attention-dependent nature of perception. Recent research has indicated that our visual experience is restricted to that to which we are currently attending. Two examples of scope flaws are change blindness and inattentional blindness. Both of these flaws are regarding our failure to notice large-scale changes and events that are occurring immediately before us. In both kinds of cases, we are unaware that we are failing to perceive any such changes and events.

3.3 FORMULATING THE PREDICTED NATURE OF EXPERIENCE

Independently of whether the advocates of the Grand Illusion hypothesis are warranted in their assumptions regarding what is indicated by the Ordinary Perceiver's surprise, let us grant them these assumptions for the sake of argument, so that we may see what these assumptions entail regarding the predicted nature of perception. To determine what is entailed, we must first formulate the predicted nature of perception with respect to each kind of flaw. We must then compare the two formulations, such that we can then see the entire account of predicted nature of experience as entailed by the assumptions of Grand Illusion hypothesis.

The Grand Illusion hypothesis appears to assume that, given what we know of the structural flaws, any visual experience which seems everywhere detailed, distinct, continuous, in color, etc., must be illusory. Suppose this were true—what does this entail about the predicted nature of perception? Since any visual experience that is “photograph-like” is illusory, and the claim that such experience is illusory is based on our expectations of how these flaws would impact our visual experience, then it follows that the predicted nature of perception would include experience of the structural flaws. For example, since our experience is said to be illusory based on the fact that we do not generally experience a lack of visual information resulting from

the lack of photoreceptors which is responsible for our blind spot, it follows that the predicted nature of experience would include such a lack of visual information.

As such, the predicted nature of perception would be highly detailed and richly colored only in a relatively small portion of the center of our visual field, with detail greatly decreasing towards the periphery. We would experience a lack of visual information in the region of the blind spot. Furthermore, our visual experience with fail to be continuous; due to the saccades, we would experience a rapid sequence of alternating black-and-white and color “snapshots”.

Regarding the scope flaws, the illusory experience is that it seems to us as though we perceive all of the changes and events that happen before us. However, as both kinds of scope flaws demonstrate, we do not notice all of them; what we perceive is constrained by that to which we are attending. Since the guiding assumption here is that it seems to us that we notice so much more of what's going on around us than we actually do, then I suppose that the predicted nature of experience with the such that it would be apparent that we only perceive that to which we are directly attending. Perhaps predicted nature of experience might be contained to some portion of the foveal region, such that the rest would be absent.

It is worth noting that here, as elsewhere, there is an equivocation on “seems”. One could mean that our visual experience is such that it seems to us as though we are noticing much more than we in fact are. This formulation seems strange to me—I am not sure what it would be like for my visual experience to seem to me as though I am experiencing much more of what is going on before me than I actually am. Maybe what is meant here is that the Ordinary Perceiver, who, for obvious reasons, does not notice that they fail to notice some of the events happen before them, forms the

implicit belief that they notice (all of the) changes going on in front of them.⁶⁰ We can then explain the surprise by appealing to that belief. However, this won't do either—the Grand Illusion is a *perceptual* illusion, not a false belief. The false belief, if not generated through illusory experience, is not indicative of a perceptual illusion.⁶¹

Note that all of these assumptions about the predicted nature of perception are derived directly from expectations generated by the research into these functional flaws. What is interesting is that the two kinds of flaws are never discussed together. They are either discussed in different sections of the same paper, or discussed in different papers altogether. As such, we can assume that the assumptions for each kind of flaw, regarding why failure to experience them is illusory, are independent. This is so because the case for each kind of flaw counting as evidence for the Grand Illusion hypothesis is made independently of any similar arguments for the other kind of flaw. Since the case for each kind of flaw is made independently, this entails that each kind of flaw has a distinct formulation of the predicted nature of perception based on the expectations generated by the assumptions regarding how that kind of flaw would affect our visual experience.

When we consider these assumptions together, our account of the predicted nature of perception becomes incoherent. The structural flaw formulation entails that we experience a lessening of color and detail towards the periphery, alternative sequences of black-and-white and color “snapshots”, no color outside the foveal region, etc., which requires the ability to perceive all of these areas of our visual field at once—which is exactly what the scope flaw formulation prevents by restricting our

⁶⁰ This is the kind of banal surprise that Jonathan Cohen attributes to the ordinary perceiver at the end of *The Grand Grand Illusion Illusion*.

⁶¹ See my earlier discussion of why this is so. Also, while it is interesting question regarding how it is that the Ordinary Perceiver has come to have such beliefs (if they do have them), such false beliefs are not themselves illusions (to call them such is a misnomer—illusions are perceptual).

experience to a portion of the foveal region. Alternately, the scope flaw formulation requires that our vision is restricted to the foveal region (or whatever region to which we are currently attending—it is generally, though not necessarily, the foveal region), which is prevented by the structural flaws, which entail that we *can* experience all of the structural flaws at once. Clearly, the result is two contrary formulations. Thus, we cannot maintain both kinds of flaws simultaneously.

3.4 ALTERNATIVES TO THE INCOHERENT FORMULATION OF THE PREDICTED NATURE OF EXPERIENCE

We have a few options. We could reject one of the kinds of flaws, establish that there are actually *two* (or more) distinct illusions, or reject the Grand Illusion hypothesis in the hope that some promising new direction might be found upon its deconstruction. Ultimately, the third option is the only viable one, because there is no principled reason for selecting one kind of flaw and rejecting the other, since they are both taken to be illusory. Furthermore, arguing that there are at least two distinct illusions will also fail because, ultimately, they are illusions of the same perceptual system, such that they will end up with the same predicted nature of perception, whether they are considered as one illusion, or separately.

It is worth being thorough in our inquiry regarding whether we could consider our failure to notice one of the kinds of flaws as an illusory experience, but not the other kind. Since it is only when we consider the together that the formulation of predicted nature of perception becomes incoherent, it seems reasonable to assume that, if we could successfully exclude one of the kinds of flaws, then formulation would no longer be problematic (at least with respect to internal coherence). In order to be able to differentiate the two kinds of flaws in this way, we need to find some theoretical grounds for arguing that the feature in virtue of which one formulation is

illusory is absent from the other formulation. This means that we must evaluate on what grounds we are assuming that each are illusions.

Let us begin with the structural flaws; recall that the expectation which motivates the claim that our failure to experience the flaws is illusory is generated by both our knowledge of the functional structure of the visual system as well as our ability to experience in some of these flaws under contrived circumstances. Thus, the claim that our experience under normal conditions (during which we fail to experience these flaws) is illusory stems from the assumption that, since these flaws are generated by the structure of our visual system, and the functional structure of our visual system is what processes our visual information, then it seems that this “hardware” would determine our experience.

Similarly, the intuition that our failure to notice the scope flaws (that attention restricts what we perceive) indicates that we are subject to an illusion rests on the failure of our visual experience to conform to the apparent constraints of the operations of our visual system. It is argued that it seems to us as though we notice so much more than we do about what is going on before us, and it is in this sense that our experience is said to be illusory.

Assume for the sake of argument that, if these are illusions, they are so for the reasons argued. We must compare the reasons that each are considered to be illusions in an attempt to determine whether there is any way that we could differentiate the two kinds, such that one is an illusion but the other is not. In the structural flaws formulation, we are subject to an illusion because our experience does not conform to what we would expect (given what we know about the structure of the visual system). Likewise, in the scope flaw formulation, we are subject to an illusion because our experience does not conform to our expectations (given what we know about the

attention-dependent nature of perception). Both cases share a common formulation: the perceiver is subject to an illusion because their experience does not conform to what we would expect their experience to be, given the kind of flaw under discussion. Furthermore, to deny that one is an illusion, is to argue that it cannot be an illusion on these grounds (it follows that if it is not an illusion, then it is not an illusion for these reasons, as well as any other reasons). But this would serve as a counterexample to the formulation on which the other illusion depends, such that it would also rule out the other kind as an illusion. As such, we cannot maintain that one is an illusion, but not the other.

Nor can we argue that, rather than both being parts of the Grand Illusion, they are two separate illusions. Suppose that we wish to argue that they are distinct illusions, in an attempt to avoid the problematic formulation of the predicted nature of perception when the two are considered part of the same illusion. The problem with this is that, ultimately, they are illusions of the same perceptual system, such that, if both are illusions, then the same predicted nature of perception is entailed—they problematic formulation would still arise.

It appears that our assumptions will not allow for a coherent formulation of the predicted nature of perception. As such, it appears that we ought to abandon the Grand Illusion hypothesis in search of a new set of assumptions which will allow for a viable formulation of our expectations regarding predicted nature of perception.

4. A NEW FORMULATION OF THE PREDICTED NATURE OF EXPERIENCE

4.1 CONSIDERING THE AFFECT THAT EACH KIND OF FLAW HAS ON OUR EXPERIENCE OF THE OTHER KIND OF FLAW

Suppose for the sake of argument that my deconstructive project is successful—where do we go from here? In other words, though deconstructive projects are useful in their own right, it may be fruitful to search the ruins of the old problem for insight into direction for new position. This appears to be just such a case.

Recall my argument for why we cannot coherently formulate the Grand Illusion hypothesis. I alleged that the Grand Illusion hypothesis resulted from faulty expectations regarding the predicted nature perception based on considering only isolated bits of data. When considering the data from the structural flaws, we assume that since we can experience, say, the blind spot, then our failure to do so under normal circumstances must be due to our being subject to illusion. Likewise, from the research on the scope flaws, we infer that, since the Ordinary Perceiver is surprised when she discovers that she fails to notice large, non-occluded events that are happening immediately before her, it must be the case that it seems to her as though she is receiving visual information from the entirety of her visual field (that is to say, from outside of the restricted area to which she is attending), when in fact this must be illusory, since we know she does not have any such experience of information external to the area to which she is attending through her failure to notice such events. From this, we have generated two contrary sets of expectations regarding what the predicted experience of perception must be like, such that deviation is grounds for arguing that such an experience is illusory. However, as we

saw earlier, this results in an incoherent formulation of predicted nature of perception. When we consider the two individual formulations together, we discover the two cannot be maintained simultaneously.

This failure is insightful, however; by examining what went wrong, we are able to see an alternative formulation of our expectations regarding our experience, given the functional flaws. The most important lesson here is that we must consider the data from the research on both kinds of functional flaws simultaneously in order to more accurately develop our expectations. The attention-dependent nature of perception is equally important of a consideration as the functional structure of visual system. For example, though we receive no visual information from the region which falls within the blind spot, given the attention-dependence of vision, namely we only notice that to which we are attending, which usually (though not necessarily) is that which falls within our foveal region, as well as the fact that are blind spot falls outside of that region, then we should not expect to experience it under normal conditions.⁶² Likewise, though it is true that we fail to notice large scale events that happen immediately before us, this is exactly sort of thing we should expect, given what we know about the attention-dependent nature of perception; if we are not attending to the change, we fail to notice it.

The previous incoherent formulation appears to arise due to emphasis on only a small portion of the causal sequence of perception. These two kinds of flaws are distinct with respect to their causal role in perception. In other words, there is more to consider with respect to what we experience than the means by which we process the visual information; quite literally, we do not “read it off the retinal image”. Instead,

⁶² Though it does seem reasonable to expect that we can experience it under conditions in which we are attending to that region. For example, we are able to experience our blind spot when we have closed one eye and, though we keep our gaze steady, we are attending to some image or symbol within that region.

there are other factors that determine our experience as well.⁶³

We can now begin to see a new formulation of our expectations when based on simultaneous consideration of the affects of both kinds of flaws on perceptual experience. If we note that the attention-dependent nature of perception prevents us from noticing that to which we are *not* attending, and we are (almost) always attending to the foveal region (which we *do* experience as detailed, distinct, in color, etc.), then it follows that we should expect that we would fail to notice the lack of detail, color, etc., outside the foveal region. As such, it appears as though our experience does conform to our expectations, and, thus, we are not subject to a Grand Illusion.

Once we have amended our expectations of what our experience must be like based on simultaneous consideration of both kinds of flaws, we see that our experience no longer fits the old framework for an illusion: our experience no longer fails to conform to what we would expect (given not only our knowledge about the flaws, but considerations of how they impact one another).

4.2 EXPLAINING THE SURPRISE WITHIN THE NEW ACCOUNT

However, I still owe an explanation for why it is that the Ordinary Perceiver is surprised when she experiences her flaws. First, remember that under the new formulation, visual experience conforms to our (the theorists') expectations. It does not follow from that, that the Ordinary Perceivers expectations will be the same as those which we have come to have the virtue of reevaluating the assumptions of the Grand Illusion hypothesis. She will come to form her expectations in a different way, namely through how her experience seems to her. We are now in a position to

⁶³ Note that I am *not* suggesting that there is some ultimate perceptual product—some finished object or image which is then perceived. Rather, I am arguing that there is more involved in determining what is the nature of our perception than the narrow evaluation of the flaws suggests.

understand why, though she is not subject to the Grand Illusion, her visual experience is such that we are nonetheless tempted to call it "photograph-like": perceptual experience under normal conditions is, due to the attention-dependent nature of perception, restricted to the visual information acquired through the foveal region. This region contains a high density of cones, allowing for her experience to be highly detailed and colored. Note that I am not positing any mechanism which somehow prevents her from experiencing the structural flaws; they are unhidden at all times, available to be perceived at any time to which they are attended but, due to the attention-dependent nature of perception, she doesn't notice them under normal circumstances⁶⁴.

If our newly-formed expectations regarding the predicted nature of perception are roughly correct, then her experience is of visual information acquired through the foveal region, which is everywhere detailed, distinct, continuous, and in color. Thus, anytime her vision does not seem that way to her, she is surprised, not because she was expecting the region containing the flaw to conform to a "photograph-like" experience of perception due to the fact that she previously experienced that region as such, but rather because she was not expecting it to fail to conform with her previous visual experiences (which, though not of that region, nonetheless fail to indicate that experience of the flaw regions would be any different). Even if we want to argue that, not only did she not expect that experience, but she also had expectations of what that experience would be like, all we could say is that she expects the same quality of detail and color that she normally experiences.⁶⁵

⁶⁴ During normal perception, she is attending to the foveal region, while the structural flaws which we have been discussing are located outside of the foveal region.

⁶⁵ Given the treatment of filling-in found within at least some aspects of the debate, it is worth mentioning that once our expectations have been reformulated in the proposed way, there is no longer the burden of explaining why we do not experience our flaws during normal perception. This has two potential implications for the advocate of filling-in: (a) some may abandon this sort of explanation since we no longer need to account for our failure to experience the flaws, or (b) some may accept the

4.3 CONCLUDING REMARKS

In conclusion, rather than expecting to experience the flaws under other conditions, our expectations can be modified by considering other, relevant constraints on perception. Attention, for example, is an important constraint: we perceive that to which we attend (and, surprisingly, little else). The structure of the eye is another important constraint, namely regarding the distribution of photoreceptors. How the information is interpreted (how we represent it as being) affects how it seems to us as well. Using these three examples, I will illustrate how constraints can alter our expectations.

We (at least generally) only notice that to which we are attending. This means that, when we are purposefully trying to find our blind-spot,⁶⁶ we are attending to that area of our visual field and noticing features of it in greater detail (or lack thereof) than when we are attending to other regions. However, we are normally attending to the foveal region of our vision. The blind-spot falls outside of that region, and is thus generally outside of our attention. Therefore, the blind-spot need not be altered for us to not perceive it—we may not perceive it because, like the rest of perception, we only notice that to which we are attending. Similar lines of thought have been used in theorizing about change-blindness and inattentional-blindness, and I see no reason why we cannot apply this explanation more generally.

The distribution of photoreceptors is another important constraint—the more photoreceptors, the greater the visual detail. Since there is a greater number of

new formulation, knowing that they no longer bear the strange explanatory burden of why we are able to experience our flaws under non-normal circumstance. Either way, it seems that, in accepting the new formulation, filling-in has been vindicated from what was generally taken (perhaps erroneously, since there appears to be no discussion of such a theory in its own right, but only as a means to this particular explanatory end) its previous role in the debate—accounting for why we don't experience the flaws—which lead some members of the debate to interpret filling-in as an unreasonable position.

⁶⁶ I do not mean for this to be the only case—there are other ways of discovering your blind-spot.

photoreceptors (and a great deal of those being cones) in the foveal region, which falls in the center of the visual field, then whatever falls into that region of the visual field will be experienced in color and great detail. We generally attend to that which falls within our foveal region, so everything to which we attend is generally experienced in color and great detail. Furthermore, the density of photoreceptors (especially cones) becomes less moving in the direction from the foveal region to the periphery. This lessens the detail and color of our vision. However, we do not generally notice this lessening of detail and color *because we are not attending to that region*. If we normally attend to the center of our visual field, and the center of our visual field is experienced as being in color and having good detail, we (a) do not notice that it is otherwise elsewhere in our visual field and (b) that which we are perceiving *is* experienced as detailed and in color (this is not illusory—it appears reasonable to argue that one’s experience really is that way in this region), which may lead us to infer that our vision (perhaps entirely) is that way. This inference is not subject to the illusion, because it is not the case that we are experiencing our *vision* as everywhere detailed, distinct, continuous, in color, photograph-like, etc. Rather, we experience the region of our visual field to which we normally attend (and which *is* that way) *as* being that way, and could mistakenly infer that our entire visual field is that way (since we do not generally attend to that which falls outside of it—we just turn our heads to look if we wish to see something that falls outside of the foveal region—and, due to photoreceptor distribution, which yields poor information outside of that region, we do not experience anything to disconfirm our experience until we experience the flaws). Given these modified expectations, it seems clear that we are *not* subject to the Grand Illusion.

BIBLIOGRAPHY

- Cohen, Jonathan. 'The Grand Grand Illusion Illusion.' *Journal of Consciousness Studies*. (9:5/6) pp. 141. 2002.
- Dennett, Daniel. 'How Could I Be Wrong? How Wrong Could I Be?' *Journal of Consciousness Studies*. (9:5/6) pp. 13. 2002.
- Dennett, Daniel. 'Filling In Versus Finding Out: A Ubiquitous Confusion in Cognitive Science.' <http://ase.tufts.edu/cogstud/papers/fillin.htm>, last accessed August 12, 2004.
- Dennett, Daniel. 'Seeing is Believing—or is it?' in K. Akins, ed., *Perception Vancouver Studies in Cognitive Science*, vol. 5: Oxford University Press, 1996 pp. 158-172.
- Noë, Alva. 'Is the Visual World a Grand Illusion?' *Journal of Consciousness Studies*. (9:5/6) pp. 1-12. 2002.
- Pylyshyn, Zenon. 'Mental Imagery: In Search of a Fundamental Theory.' <http://rucss.rutgers.edu/faculty/zpimagery.htm> Last accessed August 12, 2004.

SEEING BEYOND THE ILLUSION:
DISPELLING THE GRAND
ILLUSION HYPOTHESIS

Melissa A. Ebbers, M.A.

Western Michigan University, 2005

The Grand Illusion hypothesis is a new form of skepticism about the nature of our visual experience: it *seems* to us as though our vision is everywhere detailed, distinct, continuous, in color, even “photograph-like,” but it is not. This position is motivated by developments in perceptual research, which have revealed new information about the functional structure of the visual system as well as the attention-dependent nature of perception.

My project is primarily deconstructive. I argue that the Grand Illusion hypothesis rests on problematic assumptions (motivated by the results from the relevant perceptual research), which ultimately leads to an incoherent formulation of what we ought to expect our vision to be like. By challenging these assumptions, our expectations are altered in such a way that predicts that we would have exactly the visual experiences that we do. Thus, we are not subject to the Grand Illusion.

SEEING BEYOND THE ILLUSION:
DISPELLING THE GRAND
ILLUSION HYPOTHESIS

by

Melissa A. Ebbers

A Thesis
Submitted to the
Faculty of The Graduate College
in partial fulfillment of the
requirements for the
Degree of Master of Arts
Department of Philosophy

Western Michigan University
Kalamazoo, Michigan
April 2005

omit blank page

Copyright by
Melissa A. Ebbers
2005

ACKNOWLEDGMENTS

First and foremost, I would like to express my gratitude to the members of my thesis committee: Dr. John Dilworth, Dr. Timothy McGrew, Dr. Michael Pritchard, and Dr. Marc Alspector-Kelly. Their involvement in this process and genuine enthusiasm for the project contributed greatly to my success.

Secondly, I would like to thank my wonderful friends—including my classmates—and family (especially my mother and father) for their love and support during this process. Thank you for all of the fun and interesting discussions, as well as your kind words of encouragement.

Lastly, I would like to express great appreciation for the hard work and dedication of my advisor, Dr. John Dilworth—he certainly deserves special recognition for his contribution to this effort. His guidance and support, which included thorough comments, frequent meetings, advice, encouragement and helping me to keep a sense of humor about the project, has made this a highly rewarding experience. I am truly fortunate to have such a wonderful mentor.

Melissa A. Ebbers

TABLE OF CONTENTS

ACKNOWLEDGMENTS	ii
CHAPTER	
I. AN OVERVIEW OF THE GRAND ILLUSION HYPOTHESIS	1
1.1 Introduction	1
1.2 The Grand Illusion Hypothesis: A New Skeptical Problem?	5
1.3 How Do We Determine Whether the Ordinary Perceiver is Subject to the Grand Illusion?	5
1.4 The Role of the Ordinary Perceiver.....	11
1.5 The Ordinary Perceiver’s Belief.....	12
1.6 What Does (and Does Not) Surprise the Ordinary Perceiver	12
II. THE PREDOMINATE METHOD WITHIN THE DEBATE FOR DETERMINING WHETHER THE ORDINARY PERCEIVER IS SUBJECT TO THE GRAND ILLUSION	16
2.1 Overview	16
2.2 Distinction Regarding Methods of Belief-Acquisition.....	18
2.3 The Method of Acquiring Experience-Based Beliefs.....	21
2.4 Other Candidates for the Belief.....	24
2.5 Further Consideration of Inference-Based Beliefs	26
2.6 Why Dennett Cannot Argue That the Ordinary Perceiver is Subject to the Grand Illusion.....	31
III. REEVALUATING THE ASSUMPTIONS THAT MOTIVATE THE GRAND ILLUSION HYPOTHESIS	38
3.1 What Assumptions Motivate the Hypothesis?	38

Table of Contents

~~List of Figures~~—continued

CHAPTER

3.2 A Distinction Between Kinds of Functional Flaws.....	43
3.3 Formulating the Predicted Nature of Experience	47
3.4 Alternatives to the Incoherent Formulation of the Predicted Nature of Experience.....	50
IV.A NEW FORMULATION OF THE PREDICTED NATURE OF EXPERIENCE	53
4.1 Considering the Affect that Each Kind of Flaw Has on Our Experience of the Other Kind of Flaw	53
4.2 Explaining the Surprise Within the New Account	55
4.3 Concluding Remarks	57
BIBLIOGRAPHY	59

1. AN OVERVIEW OF THE GRAND ILLUSION HYPOTHESIS

1.1 INTRODUCTION

The Grand Illusion is an illusion of which the layperson with normal vision (let us call her the Ordinary Perceiver) purportedly suffers. As characterized by Alva Noë, the Grand Illusion is a new kind of skepticism regarding how our visual experience seems to us. The advocates of this illusion argue that it seems to the Ordinary Perceiver that her vision is everywhere detailed, distinct, continuous, in color, photograph-like, etc.,¹ when her vision is, in at least some regions of her visual field, in fact poor, not detailed and even colorless!² As a means of determining whether such a disparity exists, discussion has focused around how to interpret the surprise—and lack of surprise—exhibited³ by the Ordinary Perceiver upon experiencing her functional flaws (blind-spot, change-blindness, inattentional-blindness, saccades, etc.) as well as during normal perception.⁴ This surprise is taken to be an indication that a belief is overthrown—that the Ordinary Perceiver was expecting something else—when she discovers her functional flaws. In what follows,

¹ I will also refer to this list of attributes as “photograph-like” (as an abbreviation of the list) This list of attributes is intended to be similar to the “snap-shot” conception addressed by Noë (*Is the Visual World a Grand Illusion?*, p.2)

² This assumption, that the predicted nature of vision is not everywhere detailed, distinct, continuous, in color, photograph-like, etc., is based on a further assumption that experiencing the functional flaws is the what we ought to expect that our visual experience would be like. I challenge this assumption in Chapters 3 and 4.

³ As discussed in §1.3, I am only arguing against the predominant method (in the literature on the debate) of determining whether the Ordinary Perceiver is subject to the Grand Illusion—through interpretation of the significance of the surprise, and lack thereof, exhibited by her under difference circumstances that are relevant to the debate.

⁴ Noë argues that it is also important to note that the Ordinary Perceiver is *not* surprised by some of the adjustments (putting on her glasses, relocating, squinting, looking around, etc.) she must make to in order to perceive the desired information during normal perception. I address this in greater detail below.

I argue that the Grand Illusion hypothesis is motivated by expectations of what our visual experience must be like which are formed without consideration of sufficient relevant factors. Since these expectations are not met during our normal⁵ perceptual experience, as taken to be particularly evident given the surprise exhibited by the Ordinary Perceiver when she experiences any of her functional flaws, we assume that we are subject to an illusion. However, the surprise alone is insufficient to settle the matter since more than one belief could cause the surprise. To illustrate this point, let us discuss what motivates us to call the Ordinary Perceiver's experience illusory, namely expectations derived from recent research of the flaws.

Since the Grand Illusion hypothesis rests on expectations of what our perception must be like which are developed without considering all of the relevant factors, it is useful to attempt to formulate the predicted⁶ nature of perception based on what the Grand Illusion hypothesis takes to be illusory about the Ordinary Perceiver's perceptual experience as a means of drawing attention to the problematic assumptions which generate these expectations. My account, then, is two-fold: the first half is dedicated to what can be said regarding how the Ordinary Perceiver's visual experience seems to her, and the second discusses the implications of arguing that the Ordinary Perceiver's experience, as formulated by the Grand Illusion hypothesis, is illusory, concluding that the Grand Illusion hypothesis is motivated by unwarranted expectations (based on considering too little of the relevant data discovered through perceptual research), and arguing that once we generate our expectations by considering the affect of each of the kinds of flaw on our experience of the other, we can then see that we are not subject to an illusion—in fact, what we

⁵ By 'normal,' I mean how our experience seems to us during most, if not all, of our experiences—all but those occasions under which we perceive our structural flaws (which we do by means of non-standard usage of our perception)

⁶ See discussion of this term in §1.3.

experience is just what we would expect to experience, once we consider enough of the relevant (available) facts.

My inquiry during the first half of my project primarily addresses how the Ordinary Perceiver's vision seems to be, and incorporates other aspects of the debate only insofar as they aid in this project. I introduce the skeptical problem in Chapter 1 by addressing the functional flaws (§1.3) and what does (and does not) surprise the Ordinary Perceiver (§1.6). Through my analysis of how the Ordinary Perceiver's vision seems to her in Chapters 1 and 2, I discuss at length the role of the Ordinary Perceiver (§1.4), specifying what belief the Ordinary Perceiver must have in order to be subject to the Grand Illusion (§1.5-2.3) and whether there are any other candidates for the belief that causes the surprise (§2.4-2.5). I conclude the first half, by arguing that the case cannot be made for the alleged disparity between how the Ordinary Perceiver's vision seems to her (everywhere detailed, distinct, continuous, in color, photograph-like, etc.) and the predicted nature of her vision (discontinuous, colorless and without detail in some regions, etc.), and, as such, a new account is needed to explain the Ordinary Perceiver's reaction to various conditions. Since we cannot develop an account to explain this disparity of experience, in Chapter 3 I suggest that a new approach is adopted by the debate—re-examining our assumptions in an attempt to discover whether we may alter our expectations in such a way as to eliminate the Grand Illusion hypothesis in favor of a new explanation of our perceptual experience.

The second half of my project, then, is dedicated to what the Grand Illusion hypothesis implies about the predicted nature of perception (§3.1), which, if I am right, renders the set of expectations (generated by the assumptions made about the results of the research on the functional flaws) incoherent, illuminating an internal

problem with the formulation of the hypothesis (§3.3). In order to draw out the incoherence, I make a distinction between two kinds of functional flaws based on how they are generated (§3.2). Since it is claimed that the surprise exhibited by the Ordinary Perceiver when she experiences each kind of flaw (which are distinct kinds since they are generated by separate aspects of the perceptual system), then, since they are argued separately to be indicative of the Grand Illusion hypothesis, they admit of a distinct formulation of both how they seem to the Ordinary Perceiver, as well as of the predicted nature of perception. In other words, once the distinction is made between the two kinds of flaws, we can evaluate what the claim that the Ordinary Perceiver is having an illusory experience implies for the formulation of the predicted nature of perception regarding each kind of flaw. The two formulations of the predicted nature of perception are compared in an attempt to integrate them into one unified formulation of the predicted nature of perception. However, I argue that, once we attend to both kinds of functional flaws, the formulation of the predicted nature of perception becomes incoherent (§3.3).

Though my account is primarily deconstructive, I conclude by arguing that direction can be found for a new formulation of our expectations and the predicted nature of perception (§4.1). Once we consider both kinds of flaws (and the effect that they will have on our overall perceptual experience), we are able to formulate one account of the predicted nature of perception (rather than two separate formulations—one for each kind of flaw, and is based solely on the expectations that we would generate by virtue of that which we know about that flaw independently of how any other feature(s) of perception might effect it—which cannot be integrated into a single account). The expectations generated by the new formulation of the predicted nature of perception are such that we need no longer entertain the possibility that we are

subject to an illusion—this account entails that we should not expect to experience our flaws during normal perception (§4.3).

1.2 THE GRAND ILLUSION: A NEW SKEPTICAL PROBLEM?

Noë argues that the Grand Illusion hypothesis is a new skeptical problem: skepticism about our perceptual experience.⁷ The worry is whether our perceptual experience is different than it seems to us. The alleged disparity is that our vision seems to us to be everywhere detailed, distinct, continuous, in color, photograph-like, etc., but this experience does not accord with the predicted nature of vision, which is not everywhere detailed, distinct, continuous, in color, photograph-like, etc., due to the impairing effects of the functional flaws on our visual system. This skepticism is motivated by the surprise exhibited by the Ordinary Perceiver when she discovers one of her functional flaws. Since she is surprised, we infer that a belief was overturned—a belief, the contents of which lead her to expect that her vision would seem to her as it always had: detailed, distinct, continuous, in color, photograph-like, etc. In other words, the surprise at discovering one of her functional flaws indicates that her belief about the photograph-like nature of her vision was overturned at that time. Thus, we worry whether we ought to be skeptical about our visual experiences, given the alleged disparity between the predicted nature of our visual experiences and how they seem to us.

1.3 HOW DO WE DETERMINE WHETHER THE ORDINARY PERCEIVER IS SUBJECT TO THE GRAND ILLUSION?

Throughout the present critique of the Grand Illusion hypothesis, I restrict the scope of my evaluation to the predominate method within the debate for determining

⁷ Noë, *Is the Visual World a Grand Illusion?*

whether the Ordinary Perceiver is subject to the Grand Illusion.⁸ As such, I take issue with using the surprise exhibited by the Ordinary Perceiver when she first discovers her flaws as indication that she is subject to the Grand Illusion. In order to analyze the Grand Illusion hypothesis in the proposed fashion, we must first introduce some terms we will be using to discuss the debate. Those terms are used to must lay out the problem. Once we have clearly formulated the problem, we will be in a position to critically review the it.

Since it is my contention that the Grand Illusion hypothesis is motivated by a disparity between our expectations regarding what we would expect (based on the recent results from perceptual research), I will define my terms to reflect this. I propose the term ‘predicted nature of experience’ to refer to what we would expect our experiences to be like (relative to some set of data from which we form our expectations). Whatever assumptions regarding what we expect that our experience would be like (that we make based on the data that we are considering) are collectively considered to be the predicted nature of experience.

In contrast, an illusion (as the term is, arguably, used within the debate) is any experience which deviates from the respective predicted nature of experience. Though this is a non-standard usage of the term, something like this appears to be the implicit use of illusion within this debate. For example, given the data regarding the functional structure of the visual system (blind spot, photoreceptor distribution, etc.), we expect that our experience would be missing visual information in the region of the blind spot and would noticeably lessen in detail and color towards the periphery; these expectations comprise the predicted nature of experience. By virtue of the apparent failure of our normal experience to conform to our expectations, such

⁸ Thus, I am not speaking to other ways of arguing for the problem. I am only concerned with the general way in which the hypothesis is taken to be established.

experience is taken to be illusory.

The **functional flaws**⁹ of our visual system are aspects of the functional structure of our eyes that lead us to expect that our vision would not be everywhere detailed, distinct, continuous, in color, photograph-like, etc. The assumption appears to be that, given the existence of these flaws, we would expect that our vision is experienced as only having rich color and detail in the foveal¹⁰ region, lessening in color and detail towards the periphery, rapid sequences of snapshots, etc. Despite the presence of these flaws as part of the functional structure of our eyes, we do not notice them during normal perception¹¹ and are surprised to discover them.¹²

I am only addressing some of the flaws commonly found in the literature on the Grand Illusion that are relevant to my discussion of the disparity between how our experiences seem to us and the predicted nature of perception, such that only those flaws which can be experienced are to be taken into account of this debate. Saccades (a consequence of the function of photoreceptors), for instance, are discussed in the literature, but since I am interested in explaining the surprise exhibited by the Ordinary Perceiver when she experiences her flaws (and I am unaware of circumstances under which she could experience her saccades), I will not be discussing them.

One of the most popular of these functional flaws, the blind-spot, is the source of much excitement (and the cause of much speculation) for the visual theorist. The

⁹ I owe this term to discussions with John Dilworth. It is important to note that I am using the term “flaw” to connote that, in the debate, it is implied that these “flaws” are that which prevents us from having (the non-illusory) experience of vision as everywhere detailed, distinct, continuous, in color, photograph-like, etc.

¹⁰ The fovea of the eye is the region which we generally use during normal vision (see footnote 8). It contains a high number of cones (photo-receptors that enable us to have color vision), which allows us strong visual discriminatory power (good visual acuity, color, etc).

¹¹ By ‘normal perception,’ I mean one’s vision when one is not attempting to experience one of the flaws—this is the kind of vision we experience most, if not all, of the time.

¹² Under non-normal perceptual circumstances, such as experimental conditions.

blind-spot, which is caused by a region that contains no photoreceptors where the optic nerve connects to the eye and is located just 6 degrees from the center of the visual field (covering 16 degrees of the visual field), can be observed under certain circumstances. These circumstances differ from normal perception: during normal perception, we attend to that which is straight ahead of where our eyes are directed, while we must attend to that which is outside of the center of our vision in order to experience our blind-spot. Whatever objects or events that fall within our blind-spot are experienced (when we are *attending* to that region of our vision) as occluded from view. One of the reasons that the blind-spot is so interesting is that, though it is not far from the center of our visual field, we fail to experience it during normal perception.

Photoreceptor density and distribution is another functional flaw. We are virtually colorblind outside of the foveal region, and the level of perceptual detail diminishes from borders of the foveal region to the periphery. We do not notice either of these flaws during normal perception (as noted earlier, our experience seems to be highly detailed and richly colored during normal perception).

Change-blindness and inattentional-blindness are two related functional flaws. Both involve the failure to experience that to which we are not attending. In an experiment on inattentional-blindness, 'Gorillas in our Mist' (Simons, D. and Chabris, C., *Perception*, 1999, **28**), subjects were asked to watch a video recording of a basketball game, attending to the number of passes made by one team. During the game, a person in a gorilla costume walks onto the court, pauses to beat his chest, then exits the court. Though many of the participants correctly counted the number of passes, fewer than half noticed the gorilla. The participants were shocked when they watched the video again, observing the gorilla which they had previously not noticed.

Surprising results are also obtained through research on change-blindness.

Subjects fail to notice changes in stimuli that occur immediately in front of them. One such example is stimuli with “mud-splashes”: the subjects observe a picture, which changes every time a “mud-splash” (a shape that occludes part of the picture briefly) flashes.¹³ Subjects are surprised to learn that they failed to notice changes that occurred in the picture as they were looking at it. This sense of surprise—common to the discovery of change-blindness, inattentional-blindness and the blind-spot (as well as others)—acts as a basic step in arguing for the Grand Illusion.

Now that we have discussed the terms, we are in a position to lay out the problem. Recall that I argue that the Grand Illusion hypothesis is motivated by a deviations from the predicted nature of experience (as formulated based on the results from the research on the functional flaws). Thus, in order to determine whether she is subject to the Grand Illusion, we must answer the following questions:

- (1) Is there a discrepancy between the predicted nature of our experience and how our experience seems to us?
- (2) What is the predicted nature of our perceptual experience?
- (3) How does our perceptual experience *seem* to us?

The main question of the debate is (1). If the answer to (1) is no, then the Grand Illusion hypothesis (at least insofar as it accords with our formulation) is false; otherwise, further analysis is warranted. The answer to (1) will become more apparent in light of the answers to (2) and (3). Question (2) asks about the predicted nature of experience. According to the Grand Illusion, our vision, due to its functional flaws, is not everywhere photograph-like. In answering (2), we must address what facts about us (and the world) determine the predicted nature of vision, such that we are justified in asserting that our expectations regarding the predicted nature of perception are such

¹³ J. Kevin O'Regan offers an excellent example of change-blindness experiments that you can try for yourself at his website (<http://nivea.psycho.univ-paris5.fr/>).

that, when contrasted against the Ordinary Perceiver's experience (allegedly, that is it photograph-like), we are warranted in saying that she is subject to an illusion. Question (3) asks how experience really seems to the Ordinary Perceiver. The Grand Illusion hypothesis alleges that vision seems to the Ordinary Perceiver as though it is everywhere photograph-like. For present purposes, we must inquire how vision seems to the Ordinary Perceiver—specifically, whether it seems to her as though her vision is everywhere photograph-like. We may discover that there is no disparity between how our experience seems to us and the predicted nature of vision, in which case (1) is false. The only circumstances under which perceivers are subject to the Grand Illusion are if the answers to (2) and (3) pick out unique experiences/events/states/phenomena.¹⁴ If this is the case (if (1) is true), then we are justified in our skepticism at least to the extent that further analysis is reasonable. However, one further condition must be satisfied. Not only must there be a discrepancy between the answers to (2) and (3), but it must be the discrepancy specified by the Grand Illusion: it must seem to perceivers as though their vision is everywhere photograph-like, even though this is not the case.

Since we are concerned as to whether the Ordinary Perceiver is subject to the Grand Illusion, we need only inquire as to whether her vision seems to her to be everywhere photograph-like. If her vision does seem to her to be this way, then we are at least warranted in further investigation regarding the Grand Illusion.¹⁵ If it does not seem this way to her, then we have answered our question (at least regarding the Grand Illusion as formulated in this paper—which captures the skepticism discussed by Noë)—the Ordinary Perceiver does not suffer from the Grand Illusion.

¹⁴ I am using neutral language here to be inclusive to different theories of mind.

¹⁵ We would still need to determine whether the predicted nature of nature of experience is as it is claimed to be (in the formulation of the Grand Illusion).

1.4 THE ROLE OF THE ORDINARY PERCEIVER

One might wonder what are the conditions under which an individual may be counted as an Ordinary Perceiver, and why she is important to the debate. Surely her perceptual capacity is not so different from our own—what warrants preferential treatment for *her* beliefs? The Ordinary Perceiver is an individual who is not acquainted with visual theory. She is not aware that the structure of her visual system is such that it includes the functional flaws (mentioned above), and she is not part of the community that theorizes¹⁶ about perception. Dennett argues¹⁷ that she possesses “pre-theoretical” views about the nature of perception, based solely on how her experience seems to her. Let us grant for the sake of argument that those who argue as such are right—that she does have such commitments. These pre-theoretical views would be valuable, presumably, because they would not be biased by knowledge of visual theory, thus serving as an unbiased means for determining how experience “really” seems. Her role, then, is to shed light on whether there is a disparity (of the sort relevant to the Grand Illusion) between how our visual perception seems and the predicted nature of vision by informing us on how her vision seems to her. Thus, we need access to her pre-theoretical views to discover how vision really seems (without the influence of theoretical bias). This is, of course, a questionable assumption; it entails that we take her perception to be entirely void of theoretical bias, including bias from *folk theorizing*.¹⁸ As Dennett notes,¹⁹ there is strong evidence in support of

¹⁶ By theorizing (here), I mean the kind that the visual theorist is concerned with—she may be part of a community that engages in folk theorizing about perception.

¹⁷ *Surprise, Surprise*

¹⁸ I have benefited from discussions with John Dilworth on this matter. In my thesis, I discuss, at greater length, the possibility of an Ordinary Perceiver. I conclude that such an unbiased perceiver is not possible—at least *some* theoretical commitments are necessary in order to “know how to” maneuver oneself in order to achieve desired ends based on the information obtained through one’s perception (the fundamental thesis of the enactive approach).

¹⁹ Dennett, *How Could I Be Wrong? How Wrong Could I Be?*

the conclusion that the layperson *does* engage in folk theorizing. If she did not, then we would expect for her to not have *any* views about the nature of her own perception. But she does seem to have such views; what else could account for her confidence in her own phenomenology and her shock at the mere suggestion that it could be otherwise?

Nonetheless, for the sake of argument, let us grant that there exist Ordinary Perceivers of the variety that the debate favors—completely without commitments to the theorizing with which the Visual Theorist concerns herself²⁰—and that a perceiver without *these* commitments is sufficient for the proposed role of the Ordinary Perceiver in resolving the matter. Thus, we assume that the impact of the folk theory is insufficient to bias the experience of the Ordinary Perceiver in such a way as to destroy her function in determining whether she suffers from the Grand Illusion.

1.5 THE ORDINARY PERCEIVER'S BELIEF

We must figure out why she is surprised to discover her functional flaws—what belief was overthrown, what expectation unmet? Since the hypothesis is that it seems to the Ordinary Perceiver that her vision is everywhere photograph-like, her belief must also have the content that her vision seems “photograph-like” (and her behavior must be governed by expectations caused by this belief) in order for us to be warranted in claiming that she is subject to the Grand Illusion.²¹ This must be the belief that is overturned in order for her to have been subject to the Grand Illusion.

1.6 WHAT DOES (AND DOES NOT) SURPRISE THE ORDINARY PERCEIVER

²⁰ That is to say, at least without theoretical commitments to the functional flaws.

²¹ Due to the length of this paper, I omit my discussion of other potential explanations of why the Ordinary Perceiver is surprised by experiencing her functional flaws, such as other candidates for the belief (other than the belief argued for by Cohen).

Surprise is taken to be the indication by which we can determine how the Ordinary Perceiver's experience seems to her: it serves as empirical evidence that her experience does not accord with her expectations (which are generated by her beliefs). Since we must discover how her experience seems to her in order to determine whether she is subject to a Grand Illusion, it is crucial to address what does (and does not) surprise her.

This is particularly of interest to us, since it seems to be the prevalent method, within the debate, of attempting to establish that the Ordinary Perceiver is subject to the Grand Illusion. If we are unable to establish the Grand Illusion in this way, then those who argue in favor of it must find a new way of grounding their position. The interesting question then becomes one of how one could otherwise establish the disparity between how the Ordinary Perceiver's experience seems and how it actually is.²² Any new method of doing so must not allow any theoretical bias to enter into her interpretation of her experience—this rules out options such as asking the Ordinary Perceiver about her experience, because the language used in the questioning might bias her interpretation of how her experience seems to her.

As noted, Noë states that Dennett is one of the main proponents of the Grand Illusion and an advocate of the position that the Ordinary Perceiver's surprise is indicative of the disparity between the predicted nature of experience and how it seems to her.²³ Indeed, Dennett does appear to endorse such a view.²⁴ He argues that the surprise exhibited by the Ordinary Perceiver indicates that a belief has been overturned—that she was expecting something else. Since her surprise was caused by discovering her functional flaws, Dennett argues that she must have had a belief that

²² By which I mean the predicted nature of experience, which is assumed to be real nature of vision.

²³ Noë, *Is the Visual World a Grand Illusion?*

²⁴ Most notably in the following articles: *Seeing is Believing*; *Surprise, Surprise*; and *Filling In versus Finding Out*

led her to expect something else would occur: namely, the belief that her vision is everywhere detailed, distinct, continuous, in color, photograph-like, etc., which led her to expect that her vision would seem that way in this case, as it had seemed to her to be in all her past experiences.

In ‘Surprise, Surprise,’ a reply to Noë and O’Regan regarding the Grand Illusion, Dennett notes that his position disagrees with that of Noë and O’Regan with respect to the beliefs held by the Ordinary Perceiver. Noë and O’Regan argue that:

“But is it true that normal perceivers think of their visual fields this way [as in sharp detail and uniform focus from the center out to the periphery]? Do normal perceivers really make this error? We think not. . . . normal perceivers do not have ideological commitments concerning the resolution of the visual field. Rather, they take the world to be solid, dense, detailed and present and they take themselves to be embedded in and thus to have access to the world.”²⁵

Dennett argues to the contrary, stating that the Ordinary Perceiver *does* hold a belief that her visual experience is everywhere photograph-like. He states that such a belief can be exposed by observing the surprise exhibited by an Ordinary Perceiver during an experiment in which she discovers, say, her blind spot. He writes:

“Surprise is a wonderful dependent variable, and should be used more often in experiments; it is easy to measure and is a telling betrayal of the subject’s having expected something else. These expectations are, indeed, an overshooting of the proper expectations of a normally embedded perceiver-agent; people shouldn’t have these expectations, but they do. People are shocked, incredulous, dismayed; they often laugh and shriek when I demonstrate the effects to them for the first time. These behavioral responses are themselves data in good standing, and in

²⁵ Noë, *Is the Visual World a Grand Illusion?*

need of an explanation. They are also, of course, highly reliable signs of their “ideological commitments”—the very commitments that elsewhere in their essay the authors correctly cite as culprits that help explain resistance to their view... Surprise is only possible when it upsets belief.”²⁶

Dennett, as noted above, argues that the surprise exhibited by the Ordinary Perceiver indicates that her experience leads her to expect something other than the experience generated by the flaws. Thus, she is subject to the Grand Illusion.

Noë notes that, though the surprise discussed by Dennett must be explained, so must *lack of surprise*.²⁷ He notes that the Ordinary Perceiver is not surprised when she must put on her glasses, relocate, look around, squint, etc., to get better visual information. It seems that she would be surprised at the need for these adjustments if she believes that her vision is everywhere detailed, distinct, continuous, in color, photograph-like, etc. It is not clear why such adjustments would coincide with her expectations (thus preventing the surprise) if her experience seemed to her as though she had all the detail at once, like a photograph. He writes:

“Surprise requires explanation, but so does lack of surprise. Notice that we are not surprised or in any way taken aback by our need to move eyes and head to get better glimpses of what is around us. We peer, squint, lean forward, adjust lighting, put on glasses, and we do so automatically. The fact that we are not surprised by our lack of immediate possession of detailed information about the environment shows that we don’t take ourselves to have all of that information in consciousness at once. If we were committed to the snapshot conception, wouldn’t we be surprised by the need continuously to redirect our attention to

²⁶ Dennett, *Surprise, Surprise*

²⁷ Noë, *Is the Visual World a Grand Illusion?*

the environment to inform ourselves about what is there?"²⁸

This implies that the Ordinary Perceiver is not committed to the belief that she experiences "photograph-like" vision.

Taking stock, then, we can see that the Ordinary Perceiver *is* surprised to discover her flaws and is *not* surprised by her need to move around, put on her glasses, look around, etc., to obtain the desired visual information. This means that a satisfactory account must include an explanation for both the surprise and lack thereof. We must still determine the methods needed to develop a satisfactory explanation of this behavior.

2. THE PREDOMINATE METHOD WITHIN THE DEBATE FOR DETERMINING WHETHER THE ORDINARY PERCEIVER IS SUBJECT TO THE GRAND ILLUSION

2.1 OVERVIEW

Let us start at the beginning and work through the steps towards the conclusion that the Ordinary Perceiver is subject to the Grand Illusion (in accordance with the approach most commonly utilized within the debate). We notice that the Ordinary Perceiver is surprised to discover her flaw(s). We take her surprise to indicate that a belief has been overthrown, and we wonder what the content of that overthrown belief must have been in order for her experience of the flaw to have been surprising to her. Surprise only occurs when our experience fails to conform to our expectations. Our expectations are generated by our beliefs, which are, in turn, caused by our experiences. As such, surprise is taken to be our outward display of the overthrowing of the belief that generated the expectations which were unmet in the

²⁸ Noë, *Is the Visual World a Grand Illusion?* pg. 7

surprising situation. Her surprise, then, was caused by the failure of her experience in this case to conform to her expectations, resulting in the belief responsible for the faulty expectation being overthrown.

For example, suppose that she experiences her blind spot. In this case, her experience is such that a portion of her visual field fails to have incoming visual information. She exhibits surprise during this experience. However, from this point forward in our evaluation of the significance of the surprise, there are two ways of proceeding. At this point, I will follow the manner used in the Grand Illusion hypothesis, to see how it turns out. Later, I will return to this point²⁹ in an attempt to formulate a new interpretation of the significance of the surprise.

In order for the surprise to be generated by beliefs which she acquired through past experience (such that her belief that her vision seems "photograph-like" was generated by her experience being "photograph-like"), her past visual experiences must have generally seemed to be everywhere detailed, distinct, continuous, in color, etc. Those who argue for the Grand Illusion hypothesis must give an account of how it is that she fails to experience her flaws under normal conditions, such that she is able to have the requisite "photograph-like" experience which could generate the purported belief that her vision is "photograph-like."

However, her surprise alone is still insufficient to warrant attribution of the Grand Illusion to the Ordinary Perceiver. While it can be argued that the surprise exhibited by the Ordinary Perceiver when made aware of their functional flaws for the first time is indicative of a belief being overthrown, the debate needs a means of explaining and resolving the disagreement between those who argue that this is sufficient evidence that the Ordinary Perceiver is subject to the Grand Illusion and

²⁹ Chapter 4

those who argue to the contrary. The crucial distinction, I would suggest, lies in the method of belief acquisition: how did she come to have the belief in question? To this end, I propose a distinction between two ways that the relevant belief could have been acquired: experience-based beliefs versus inference-based beliefs.

2.2 DISTINCTION REGARDING METHODS OF BELIEF-ACQUISITION

Experience-based beliefs (call them ‘belief_e’) are acquired through perceptual experience. Inference-based beliefs are acquired through inferences made from other beliefs. In other words, I acquire an inference-based belief, (call these ‘belief_i’) as a result of inferring the content of that belief_i from another belief (which may itself either be a belief_i or a belief_e). An example will be helpful to illustrate the distinction. Consider the following incident which happened to me recently while shopping for a new shirt. I was studying a particular shirt, commenting to my companion that I liked the shirt very much except for the pattern. My companion then drew my attention to another shirt across the table, and asked whether I liked that shirt better. This shirt had a pattern that I liked much more. I decided that it was better than the other one, since, other than the more preferable pattern, it seemed to be qualitatively identical in all other respects. It was only after I had made that statement aloud regarding my judgment that the two shirts were otherwise alike that my companion to my attention to an interesting detail that I had missed: the shirt was not, as I had thought, similar in all other respects—the collar was entirely different (the first had a button-less v-neck collar, while the second was a polo shirt with three buttons).

My belief that this shirt had the same kind of collar as the first was inferentially acquired. Was I subject to illusion in this case? Though I judged the two shirts to have similar collars, it does not follow that I had illusory experience. I made my judgment about the second shirt without having any experience of the collar at

all—it was not the case that I had an *illusory* experience of the second shirt as having the same collar as the first (such that I would have acquired a belief_e that the collars were the same style); instead, that the collar seemed to be the same was a judgment that was not based on perceptual experience, illusory or otherwise.

The proposed distinction does important work for us regarding what is meant by ‘seems to be the case to the Ordinary Perceiver.’ Discussion of how something *seems* to an individual appears (at least in this debate) to refer to the *experience(s)* that are being had by the individual. The distinction discerns between cases where it *seems* to me (experientially) that the second shirt has the same kind of collar as the first, and cases where I infer that the second collar is the same as the first (perhaps I implicitly assumed that all of the shirts on that table must have the same kind of collar, and I failed to have any experiences of the shirt that suggested that it was not so). It is only in the experiential cases that we have the “seeming” that is required for the illusion. The beliefs_i, in contrast, only tell us that it is not the case that it seems to the Ordinary Perceiver that her vision is not everywhere detailed, distinct, continuous, in color, photograph-like, etc. This is so because she is not receiving the requisite *experiential* information to form a belief regarding a positive statement about how her experience seems to her; instead, the information that she *is* receiving (primarily from the foveal region) seems photograph-like and, given no information that it is otherwise, she *infers* that the rest of her visual field is this way as well.

It is important to note the distinction between what is meant by “seems” when it is used to describe the content of the two kinds of beliefs. Regarding the content of the experientially-acquired belief, by “seems”, we mean that, to the perceiver, the content of the perceptual experience was sufficiently robust to cause the belief. By “sufficiently robust,” I mean that the nature of the perceiver's experience is sufficient

to cause the belief (rather than relying on the kind of inference that led to my belief that the shirts has the same kind of collar). In my earlier example, the experientially-acquired belief that the second shirt had the same kind of collar as the first could only have been generated by having an (illusory) experience of the shirt as having that style of collar. In contrast, “seems”, as used to describe the content of an inferentially-acquired belief, means that the perceiver judged it to be so without having actually experienced it as such. For example, though I might have been surprised to learn that the second shirt had a different style of collar than the first, and I might have exclaimed the collars seemed to be the same to me, my belief that the second collar was the same style was not acquired through a perceptual experience of the second shirt as having a v-neck collar. Therefore, it seems incorrect to say that I was subject to an illusion—rather, I acquired the belief through the fact that the part of the shirt to which I was attending was similar in style to the first, and that I was not perceiving any information to the contrary; in fact, I had not glanced at the collar at all. The fact that I had not, at any point during which my belief about the similarity in style between the two shirts was formed, looked at the collar makes it very difficult for one to argue that I was subject to an illusion—I cannot be said to have experienced the collar of the shirt as a button-less v-neck (which would have been the illusory experience) if I did not experience the collar at all.

Thus, we can see the importance of the belief being acquired experientially. This difference between the content of these beliefs (though *functionally* equivalent)³⁰ matters because, in order to attribute to her the Grand Illusion, it must *seem* to the Ordinary Perceiver that her vision is everywhere photograph-like.³¹ This “seeming”

³⁰ By functionally equivalent, I mean that both beliefs would cause the Ordinary Perceiver to exhibit surprise when she discovers her flaws.

³¹ I suppose that it could be argued that inference-based beliefs could be sufficient to motivate the Grand Illusion hypothesis, but this seems to be a non-standard use of the term ‘illusion’. Illusions are understood to be experiential and, though the Ordinary Perceiver would exhibit the surprise if the

can only be established in the case of beliefs_e. Thus, we must determine whether this belief, if held by the Ordinary Perceiver, could have been experientially-acquired such that we are warranted in attributing the Grand Illusion to her.

2.3 THE METHOD OF ACQUIRING EXPERIENCE-BASED BELIEFS

How could she acquire the requisite belief experientially? Since she could only do so through experiences that seemed everywhere photograph-like, we must determine whether such an experience is possible. Note that there are two criteria that our account of her experience must satisfy: her experience must seem photograph-like during normal perception, but we must also account for how she is able to experience her flaws under non-normal conditions. One suggestion is that the functional flaws are “filled-in.”³² “Filling-in,” at least as it is understood with respect to this debate, refers to a theoretical process whereby the brain edits and adds information (hence “fills-in”) in the regions of the visual field that are missing visual information due to functional flaws—allegedly accounting for why we do not notice these gaps of missing information.³³ Suppose that we wish to entertain some variety of “filling-in.” This would allow for the Ordinary Perceiver to have the experience requisite for acquiring the belief. Would it also warrant us in attributing to her the Grand Illusion? It does not: her belief that her vision is everywhere detailed, distinct, continuous, in color, photograph-like, etc., would be *true*; her experience would not be illusory

experience failed to conform to her expectations as generated by an inference-based belief, it seems to be a misnomer to call her previous experience an illusion, in the same way that it seems incorrect to say that I suffered from an illusion in my “shirt collar” example.

³² I intend for this term to be understood, as used here, as referring to a theoretical process whereby the “gaps” that are a consequence of the flaws are touched up in such a way (the bit-map “filling in” of Dennett’s *Filling-in vs. Finding Out*) that the gaps are eliminated. We perceive the experience as “gap-free.”

³³ I contend that the argument that we ought to experience these “gaps of information” under normal circumstances (or, at least, that is what perception *really is like*) is invalid. This is because it does not follow from the fact that we can experience these flaws under certain circumstances that either (a) we ought to expect to experience them at all times, or (b) to not experience them under other circumstances is indicative of being subject to an illusion.

because the filling-in would make her vision seem that way because it *is* that way. In other words, according to filling-in, the perceptual product—the object of perception—is everywhere detailed, distinct, continuous, in color, etc., such that photograph-like perceptual experience is veridical. However, the challenge³⁴ for those who would argue for filling-in is to explain how it is possible for us to *ever* experience our flaws. Given that the nature of “filling-in” is such that that we do not experience the flaws because what we perceive is a filled-in perceptual product which is developed further down the causal chain (retinal image is touched up), it is not clear why filling in would fail us (since the image does not depend strictly on the incoming information—it is an edited version of the original information).

Perhaps an advocate of filling-in might object to my dismissal of her position as a viable explanation for the Ordinary Perceiver’s experience.³⁵ After all, doesn’t our failure to experience these flaws warrant *some* explanation—and what could be better than an explanation that argues that we would not expect to experience them because they are “filled-in”? We might inquire as to how the “filling-in” proceeds—that is, how it is that these flaws are edited. The advocate of a veridical variety of “filling-in” could answer, as so many have, that the brain “knows what goes there” and “fills it in” accordingly. However, this explanation has a bizarre, undesirable consequence: we would expect that, in unfamiliar environments, our functional flaws would initially be exposed until our brains figure out “what goes there.” Since we do *not* experience this occasional exposure of the flaws,³⁶ we can reject explanations that

³⁴ The need to meet such a challenge is removed once the Grand Illusion hypothesis is rejected in favor of my proposed reformulation of our expectations. See footnote 29, below.

³⁵ See my argument for how, once we have modified our expectations appropriately, the advocate of filling-in no longer owes any explanation regarding why filling-in “fails” when we perceive our flaws. Furthermore, all of the undesirable consequences (both here and elsewhere in the debate) are removed, such that filling-in can be seen as a viable theory of perception, rather than the oddity that it is taken to be at this point in the discussion.

³⁶ I can assert this with confidence, given that Ordinary Perceivers exhibit surprise upon experiencing their flaws, a reaction that we would not expect if they experienced them occasionally; it is the novelty

posit veridical filling-in³⁷

Suppose that one were to argue for a non-veridical variety of filling-in. The bizarre consequence of the veridical filling-in account would be avoided—but at what cost? One might argue that the content need not be veridical—any filling-in will do, just so long as the perceiver does not experience the gaps. However, this has a bizarre consequence of its own. Depending how this filling-in is selected, we might be *more* likely to notice our functional flaws: if the assignment of the filling-in design (what the filling-in looks like) is random, then we would expect our functional flaws to be obvious to us. The content of the filling-in could be different from the incoming visual information around it, such that we could more easily detect the flaws due to the anomalous information.³⁸ The advocate of non-veridical filling-in might counter by arguing that, since these flaws are not located in the foveal region, like the rest of the visual information that we acquire external to the fovea, we do not experience it in enough detail to notice the difference. However, this claim is not valid: while their assertion about not experiencing visual information acquired external to the foveal region as highly detailed is true, the same argument cannot be made for the contents of filling-in—unlike the visual information, which requires photoreceptors to be acquired, filling-in is argued to be generated by the brain, which is not subject to the same restrictions. One could, of course, argue that the brain represents that content as poorly as the visual information attained through that region external to the fovea (perhaps so as to not create disparities between the “filling-in” and the surrounding visual information), but this begins to look like an unnecessarily complex

of the experience which warrants the surprise (and which would be lacking if this explanation were true).

³⁷ At least insofar as the account of “filling-in” requires “knowing what goes there” (and, perhaps due to my unfortunate lack of imagination, I cannot fathom any other way of attaining veridical “filled-in” content).

³⁸ Rather than experience the flaws as poor or absent visual information, we would now experience them as containing anomalous information as compared to the rest of the visual field.

explanation.

Could our Ordinary Perceiver gain the experientially-based belief any other way? It seems that she could not. Her being able to acquire the belief through experiencing her vision as everywhere detailed, distinct, continuous, in color, photograph-like, etc., requires that her vision *is* this way (thus enabling her to experience it as such).³⁹ However, if her vision *is* this way, then she is not subject to any illusion—the content of her experience does not deviate from the predicted nature of experience⁴⁰ and her belief that her experience is “photograph-like” is true.⁴¹ Indeed, any experientially-based belief that is acquired through experiencing veridical, filled-in content will render (1) false; there is no discrepancy between the nature of our vision and how it seems to us. This rules out the experientially-based belief with the content ‘my vision is everywhere detailed, distinct, continuous, in color, etc.’ as a viable option.

2.4 OTHER CANDIDATES FOR THE BELIEF

In order to be justified in attributing the Grand Illusion to the Ordinary Perceiver, we must discover whether the belief that her vision is everywhere

³⁹ There is, perhaps, one other option, which I do not entertain here (because I do not see any reason to take it seriously): continuous hallucination. This would not require “filling-in” (this is different from the brain “touching-up” the contents, and could be entirely accidental) but would involve constant hallucination in the regions of those flaws.

⁴⁰ This is because the object of perception is, on such an account of filling-in, a touched-up image, which *is* everywhere detailed, distinct, continuous, in color, etc. Given that an account that uses filling-in as an explanatory resource must posit that perception is edited in this way, then our expectations would be altered from the original set in the debate—we would expect that our vision is everywhere detailed, distinct, continuous, in color, etc., because that which we perceive (the filled-in perceptual product) is that way.

⁴¹ Further, it is worth noting that it is unclear how we could *ever* experience the flaws (if they are “filled-in”). Perhaps the guise is only active when we are not attending to it or, like a poorly placed mask (think of enjoying a play from a seat where one’s angle to the stage is such that one only sees the sides of the masked actors—we would see what the person on stage *really* looked like (behind the mask), and not the mask—whereas the rest of the theatre enjoys the disguise), we are unable to see it under certain circumstances. Still, this begins to feel ad hoc—the more “just so” our account becomes, the more we may rightly long for a stronger explanation.

photograph-like, is the only belief that would cause her to exhibit such surprise and, if not, whether it is at least equally as strong of an explanation as the attribution other beliefs (regarding how well it fits the relevant information). As such, we must evaluate other candidates for the belief that would cause her to exhibit this surprise. If any of these candidates are viable, we must determine whether the belief_e that her vision seems photograph-like has at least comparable, if not greater, explanatory strength.

To illustrate that viable alternatives exist,⁴² we will address one such alternative raised by Jonathan Cohen.⁴³ He suggests that, rather than having theoretical commitments regarding the nature of her visual field (which he believes that we are not warranted in attributing to her), the Ordinary Perceiver simply believes that they notice objects and events that are located or occurring in front of them. Such a belief would also provoke surprise in light of experiencing the functional flaws—it would generate expectations for the Ordinary Perceiver that (it seems to her as though) she never fails to notice objects and events that are directly in front of her. The Ordinary Perceiver would be surprised upon discovering her blind-spot (where part of the object(s) are occluded), experience change-blindness or inattentional-blindness (failing to notice events happening before them due to lack of attention) because these experiences fail to conform to the expectations that would be caused by such a belief.

As such, the requisite belief is not the only one which would elicit surprise in

⁴² Due to the length of this paper, I omit my discussion of other potential explanations of why the Ordinary Perceiver is surprised by experiencing her functional flaws, such as other candidates for the belief. The importance of illustrating that there is at least one other candidate for the belief serves as evidence that we must continue our inquiry; that the Ordinary Perceiver is surprised—and that the surprise *can* be explained by the requisite belief—is insufficient grounds for attributing to her the Grand Illusion. However, see also my later discussion of beliefs about the nature of occlusion as another potential explanation.

⁴³ Cohen, *The Grand Grand Illusion Illusion*

the Ordinary Perceiver. Since this is so, we must further examine the beliefs that we could attribute to the Ordinary Perceiver to determine which is the best explanation of her surprise (and whether the surprise is indicative of the Grand Illusion).

2.5 FURTHER CONSIDERATION OF INFERENCE-BASED BELIEFS

If I am correct in concluding thus,⁴⁴ then the only candidate (with the content ‘my vision is everywhere detailed, distinct, continuous, in color, etc.’) for the cause of the Ordinary Perceiver’s surprise is a belief_i. This belief_i would cause her to exhibit surprise upon discovering the functional flaws of her visual system. The Ordinary Perceiver would have the tacit belief⁴⁵ that it is not the case that her vision seems to her as though it is not everywhere detailed, distinct, continuous, in color, photograph-like, etc. One might wonder why I assert that beliefs_i are held *tacitly*. Arguably, not all our beliefs are held overtly; some of our beliefs, though we may not actively entertain them or have the ability to express them linguistically (or even know that we have these beliefs), nonetheless impact our behavior. For example, though an Ordinary Speaker of the English language may not have explicit access to her beliefs about the conjugations of words in the English language (and, given what I know about the teaching of the English language in most schools, it is safe to assume that most Ordinary Speakers of English do not), she can nonetheless detect the occurrence of a conjugation error. This is similar to what Pylyshyn refers to as ‘tacit knowledge’: though the individual might not be able to make use of such information for the purposes of answering questions about the matter, she can reveal such knowledge by

⁴⁴ That there is no viable account that would explain how she could experientially-acquire the requisite belief while nonetheless still be able to experience her flaws under non-normal perceptual circumstances.

⁴⁵ By ‘tacit belief,’ I mean roughly what Pylyshyn describes as ‘tacit knowledge,’ though I do not intend to imply any further similarities between the ‘mental imagery debate’ and the Grand Illusion at this point.

exhibiting telling behavior when occurrences are contrary to her knowledge.⁴⁶ Likewise, this would account for both the surprise exhibited by the Ordinary Perceiver upon experiencing the flaws *and* her strange looks and responses when queried about whether she believes that their vision is everywhere detailed, distinct, continuous, in color, photograph-like, etc.⁴⁷

Perhaps one might object that my explanation is incomplete yet; I have not yet shown how inference-based beliefs do not warrant our attribution of the Grand Illusion to the Ordinary Perceiver. After all, it appears that the content of the inference-based belief is at odds with the “real” nature of perception.

To explain how inference-based beliefs do not warrant the attribution of the Grand Illusion to the Ordinary Perceiver, it is useful explicating two seemingly disparate accounts by Noë and Dennett, illustrating how these accounts can be interpreted as sharing much common ground. This common ground illustrates that the surprise, if generated by inference-based beliefs, is insufficient to warrant attribution of the Grand Illusion to the Ordinary Perceiver.

In an attempt to explain why the Ordinary Perceiver is surprised by her functional flaws, Noë introduces the problem of *perceptual presence*.⁴⁸ This problem is a strengthened version of the Grand Illusion hypothesis. The worry is that we seem to “experience” a whole object when we are, in fact, only perceiving part of it. For example, when we close our eyes and hold a bottle in our hand, it seems to us as though we are experiencing the whole bottle when, in fact, our hand is only making contact with a portion of it. Likewise, when we see a cat on the other side of a fence, though it seems to us that we experience the whole cat, we only perceive part of it

⁴⁶ Pylyshyn, *Mental Imagery: In Search of a Theory* (particularly section 3.1)

⁴⁷ Ordinary Perceivers generally reply to inquiries about their visual experience with either puzzled looks or thoughtful replies stating that either they do not believe their experience to be that way or that they had never thought about it.

⁴⁸ Noë, *Is the World a Grand Illusion?*

(that which is visible between the boards of the fence).

Noë argues that the key to resolving the problem of perceptual presence is through analysis of amodal perception. Amodal perception is, “perceiving what is, strictly speaking, out of view.”⁴⁹ Just as we amodally perceive the hidden parts of the cat, Noë suggests that we amodally perceive objects in our environment which we are not currently *visually* experiencing. He states that this proposed alteration to the problem changes the question from one of whether the Ordinary Perceiver has the (illusory) experience as though she perceives everything in the environment at once to one of whether amodal experience is illusory. Noë argues that there is no need for the brain to create a representation of, say, the whole cat (which is partially occluded by the fence) since all we need to do to get more information about the cat is to move around so that different parts of the cat become visible. We are skilled perceivers; we know how to manipulate objects and adjust ourselves such that we are able to perceive the information that we seek. When we squint or peer, as Noë puts it, we do so because that which we wish to see is yet too far or too small from our current visual perspective. We know what our range of visual clarity is and we adjust accordingly. Too big? Back up. Too small? Lean in.

At the end of the day, though Noë and Dennett have argued against one another’s positions regarding whether the Ordinary Perceiver is subject to the Grand Illusion, it appears that they have much in common. Both argue that the Ordinary Perceiver does not *experience*⁵⁰ vision that seems everywhere detailed, distinct, continuous, in color, photograph-like, etc.; rather, their experience is a matter of

⁴⁹ Noë, *Is the World a Grand Illusion?*

⁵⁰ Though Dennett claims that it does *seem* that way to the Ordinary Perceiver—contrast with Noë, who argues that it does not seem that way to her (the phenomenology involved in the Grand Illusion hypothesis is in error with respect to how it seems to her). Noë argues that, rather than perceiving a line which extends through the region of your blind spot (while you are attending to that region) as being unbroken, it seems to her as through her experience of the line is that it does not seem to be broken

representation, interpretation, or skilled (amodal) perception of those regions that fall outside of the region to which the perceiver is currently visually attending. Both accounts rely on a judgment sense of “seeming” (amodal perception, “finding out”), not on the experiential sense of the term. Understood in this way, both of their positions are consistent with arguing that the relevant belief held by the Ordinary Perceiver is an inferentially-based⁵¹ belief. The ‘amodal perception explanation’ offered by Noë appears to be very similar to the account that Dennett gives as an explanation of our experience of Andy Warhol’s Marilyn Monroe wallpaper—both argue that we do not *experientially* represent the poor or absent information as being present. Dennett writes:

“Suppose you walk into a room and notice that the wallpaper is a regular array of hundreds of identical sailboats, or—let’s pay homage to Andy—portraits of Marilyn Monroe. We know that you don’t foveate, and don’t have to foveate, each of the identical images in order to see the wallpaper as hundreds of identical images of Marilyn Monroe. Your foveal vision identifies one or a few of these and somehow your visual system just generalizes—arrives at the conclusion that the rest is “more of the same.” We know that the images of Marilyn that never get examined by foveal vision cannot be identified by parafoveal vision—it simply lacks the resolution to distinguish Marilyn from various Marilyn-shaped blobs. Nevertheless, what you see is not wallpaper of Marilyn in the middle surrounded by various indistinct Marilyn-shaped blobs; what you see is wallpaper composed of hundreds of identical Marilyns. Now it is

⁵¹ Presumably, Noë would not use the term “inference.” However, I only intend to imply that both his account and Dennett’s “finding out” involve something like a judgment. For example, on both accounts the belief that vision seems “photograph-like” is a judgment about the nature of vision based on the parts of the visual field that the perceiver can experience (and, from that, she judges that the flaws contain similar visual information), and not because she *experiences* the flaws as being “filled-in”, in the sense that there is some perceptual product which she *experiences* as “photograph-like”.

a virtual certainty that nowhere in the brain is there a representation of the wall that has high-resolution bit-maps that reproduce, xerox-wise, the high-resolution image of Marilyn that you have foveated. The brain certainly would not go to the trouble of doing that filling in! Having identified a single Marilyn, and having received no information to the effect that the other blobs are not Marylins, it jumps to the conclusion that the rest are Marylins, and labels the whole region "more Marylins" without any further rendering of Marilyn at all. Of course it does not seem that way to you. It seems to you as if you are actually seeing hundreds of identical Marylins. And in one sense you are: there are, as I said, hundreds of identical Marylins out there on the wall, and you're seeing them. What is not the case, however, is that there are hundreds of identical Marylins represented in your brain. Your brain just represents that there are hundreds of identical Marylins."⁵²

Notice another similarity between their accounts: they both offer interpretations that fit the behavior of the Ordinary Perceiver. Both accounts are consistent with the surprise (due to experiences that are contrary to her expectations) and lack thereof (due to experiences that conform to her expectations) that is exhibited by the Ordinary Perceiver. Neither account attributes to her expectations (or beliefs) that would lead her to fail to be surprised when experiencing the flaws—her surprise is consistent with both accounts, and the same is true for her lack of surprise. Thus, the assertions of each regarding what does and does not surprise the Ordinary Perceiver fit equally as well with the other's account as their own.⁵³

I will now consider whether the content of the inference-based belief is such

⁵² Dennett, *Filling In versus Finding Out*

⁵³ This is particularly interesting when one considers that the role that Dennett's assertions regarding what *does* surprise the Ordinary Perceiver and Noë's assertions about what does *not* surprise the Ordinary Perceiver each appear to be intended as evidence against the other's account. However, in this light, it appears that both sets of assertions fit *each* account equally well.

that we are warranted in attributing the Grand Illusion to the Ordinary Perceiver. Recall that the inference-based belief in question is held tacitly. If the Ordinary Perceiver tacitly believes that it is not the case that her vision seems to her as though it is not everywhere photograph-like then her belief does not entail that she also believes that her vision *is* everywhere detailed, distinct, continuous, in color, photograph-like, etc. No beliefs that assert positive claims are made about how her vision seems to her—only a negative one: it isn't the case that her vision is not photograph-like. She could acquire this negative belief because she has no reason to accept its negation: that her vision is not everywhere photograph-like. In other words, in order for her to acquire the positive belief that her vision is not everywhere photograph-like, she must have an experience that causes her to have that belief (such as experiencing the flaws). In the absence of such experiences (and experiences that would cause the belief that her vision is everywhere photograph-like—such as those discussed above), she is committed to neither belief and instead acquires the tacit belief that it is not the case that her vision seems to her as though it is not everywhere photograph-like. This tacit belief roughly amounts to the expectation discussed by Cohen: the Ordinary Perceiver believes that she perceives the events and objects that are directly before her.⁵⁴ In other words, her expectation is governed by her tacitly held belief that it is not the case that her experience is not everywhere photograph-like (she does not expect that she will fail to experience some of the objects and events immediately before her).

2.6 WHY DENNETT CANNOT ARGUE THAT THE ORDINARY PERCEIVER IS SUBJECT TO THE GRAND ILLUSION

Recall my earlier argument that the belief which is allegedly overthrown when

⁵⁴ Cohen, *The Grand Grand Illusion Illusion*

the Ordinary Perceiver is surprised at the experience of her flaws must satisfy two criteria: (a) it must have the content that vision seems to be everywhere “photograph-like,” and (b) this belief must be acquired through experience (her vision must seem “photograph-like,” in the experiential sense of seeming). We have seen that one cannot argue for the Grand Illusion by using the surprise as evidence unless the belief in question was acquired through experience. Now we will look at why one cannot make the case for the Grand Illusion unless one argues that beliefs have determinate content.

As noted earlier, Noë argues that Dennett is the main proponent of the Grand Illusion Hypothesis. However, I contend that he cannot argue for the Grand Illusion in the way that he attempts to do so. Since he is taken to be a central figure in the debate, it is worth reviewing the problem with his account which prevent him from being able to argue for the Grand Illusion, as well as one regarding his explanation for why the Ordinary Perceiver fails to experience her flaws.

There are two problems with Dennett’s account of the Grand Illusion hypothesis: the first is regarding how he attempts to establish his case for the Grand Illusion, and the second is regarding whether his “labeling” violates his Thrifty Producer Principle. Since Dennett’s account of how he attempt to establish his case for the Grand Illusion hypothesis is similar to other accounts that we have discussed which support their claim by noting the surprise exhibited by the Ordinary Perceiver, we are already familiar with that strategy. However, Dennett’s case is special, in that his holism appears to prevent him from being able to establish the Grand Illusion hypothesis. Once we have seen why this is so, we may then inquire as to whether one aspect of his account regarding why we do not experience our flaws violates a central principle in his general position.

Dennett argues that the surprise exhibited by the Ordinary Perceiver when she first experiences her flaws is indicative of a belief being overthrown—a belief that her vision is everywhere detailed, distinct, continuous, in color, etc. However, Dennett, unlike other advocates of the Grand Illusion hypothesis, does not mean that she has a belief with such determinate content; as a holist, he argues that beliefs have indeterminate content. It seems consistent with his overall position to interpret him as attributing this belief to her as part of an intentional stance that he adopts towards the Ordinary Perceiver. As such, the belief is attributed as a means of explaining her behavior.

This seems right so far as it goes—if she had such a belief, *then* we would expect her to exhibit surprise when she first experiences her flaws. The interesting question is whether there is reason to think this is the *best* stance to adopt towards her behavior. Important considerations regarding whether this is the best explanation⁵⁵ include how well it integrates with the rest of the beliefs that we attribute to her when we have adopted stance with the strongest explanatory power. In other words, are there beliefs that we attribute to her which would have stronger explanatory power regarding her behavior, as well as fit better with the rest of the best stance that we adopt towards her?

Compare this belief to one regarding the nature of occlusion. For example, it seems that we could explain her surprise does not indicate when she discovers her blind spot by attributing to her a belief about the nature of occlusion (as supported by her experiences thus far)—namely, the belief that an object (or some part thereof) is occluded from view if and only if it is blocked by another object. Surely, to one who

⁵⁵ Since his Intentional Stance is a predictive and explanatory strategy, meaning that we adopt a particular stance both to make sense of previous behavior and as a means of predicting future behavior (by hypothesizing how an ideally rational system might behave).

holds such a belief, it would seem surprising for an object to be occluded from view without being hidden behind another object. However, such surprise does not indicate that she is subject to the Grand Illusion; though her experience in this case fails to conform to her expectations, it is this, and not an illusory experience, that generates the surprise. As such, it appears that there are other beliefs that we can attribute to her which explain the surprise but which fail to indicate that she is subject to illusion.

Regarding which belief we should attribute to her, we need to consider which best fits with the rest of the intentional stance that we adopt towards her. While the former belief fails to connect up with any other beliefs we can attribute to her (perhaps this is why we are so resistant to saying that the Ordinary Perceiver really does have such a belief—it is so unlike the other beliefs that we could attribute to her that we have a hard time attributing such a belief), it appears that the latter could easily fit within the framework of cognitive development. For instance, if we assume that she acquired her belief about the nature of occlusion through previous experience, namely by having noticed that the only time that objects disappear is when they are hidden behind other objects, then she would not expect for an object to be occluded in that region of visual field (given that there is no object there which would occlude it).

Regardless of whether we can determine which belief best fits with intentional stance that has the greatest explanatory and predictive power, if attributing a belief with this content to the Ordinary Perceiver is necessary for making the case that she is subject to the Grand Illusion, then we ultimately end up with a skeptical problem. To understand how this problem arises we must first get clearer on how, exactly, Dennett attempts to make his case. Once we see this, it will become apparent that, in light of his holism, Dennett's position entails that there is no fact of the matter whether the

Ordinary Perceiver is subject to the Grand Illusion.

Recall that the importance of belief is merely to reveal the content of the Ordinary Perceiver's visual experience. When she exhibits surprise, we can infer that her experience failed to conform to the expectations generated by her beliefs. Thus, the surprise she exhibits upon first discovering her flaws is taken to indicate that a belief with the contents which would generate expectations which are contrary to her experience of flaws, which was generated by past experience, is overthrown.

Holism, however, entails that no beliefs have determinate content; the content of beliefs are indeterminate, meaning that there is no fact of the matter what is the content of belief. This is problematic because, as we saw earlier, the belief with requisite contents is necessary for establishing the Grand Illusion hypothesis in this way. Regardless of which belief we attribute to her, Dennett must⁵⁶ conclude that there is no determinate content to her belief, such that it is indeterminate whether she is subject to such illusion.⁵⁷

Regarding the second problem with his account, Dennett's explanation for why we fail to experience our functional flaws has two parts: (1) the brain ignores the missing data and (2) the brain "finds out" what information falls within the areas of missing data, and labels that region accordingly. Central to his argument that we ought to reject filling-in in favor of his account is his Thrifty Producer Principle—"if there's no one to look at it, don't bother making it". In other words, if there is no part of the brain designated to processing such information, it is superfluous for such information to be generated (presumably by whatever mechanism would be

⁵⁶ In light of his holism.

⁵⁷ Note that I am not arguing about whether beliefs have determinate content. I am only arguing that, given the use of the surprise as the means of determining whether the Ordinary Perceiver is subject to the Grand Illusion, the belief in question must have determinate content to make the case in this way. But this, of course, is independent of my own position regarding whether beliefs are determinate, which is irrelevant here.

responsible for the filled-in information). However, this principle in conjunction with the two parts of his explanation for why we fail to notice her flaws makes finding out superfluous as well as a violation of principle.

His account begins with arguing the brain simply ignores the missing data. He maintains that there is no part of the brain dedicated to processing visual information from the regions the visual field which, due to the perceptual flaws, lack such information to some extent. Since there is no part of the brain assigned to the task, the missing information goes unnoticed.

It is this claim that there is no part of brain assigned to the task of processing what would be the visual information from the regions which contain the flaws which motivates his Thrifty Producer Principle; that there is no part of the brain assigned to task seems (to Dennett, anyway) to entail that there is no part of the brain which would perceive the touched-up perceptual product, making any such product superfluous. Therefore, he concludes that there must not be any filling-in.

He then goes on to argue that the brain “find out” what falls under those regions of the visual field which are missing information due to the flaws and labels that region accordingly. However, this part of the account is problematic for two reasons. The first is that it is superfluous, given that he has already argued that the brain ignores the missing data. Secondly, it violates the thrifty producer principle.

An example of how “finding out” works will be useful here. Recall the earlier quoted passage in which he illustrates this point through asking the reader to imagine that she is standing before a wall full of Marilyn Monroe image wall paper by Andy Warhol. He argues that, though it seems to the Ordinary Perceiver as though she is seeing hundreds of Marilyn images, is not the case that, somewhere in her brain, there is a representation of hundreds of Marilyn images; rather, from the images that she

does see before her, and no disconfirming information to the contrary from the periphery, the brain just “jumps to the conclusion” that it is more of the same, and labels that region “more Marilyn images”.

Finding out is unnecessary given that he has already argued that the brain ignores the missing data. If the missing information is ignored, then it is not clear why we need to posit the labeling process. More specifically, it is not clear what work the labeling is doing. The labeling appears to be intended to play an explanatory role in his account regarding why we do not notice the flaws, but the ignoring feature of his account appears to fully explain this independently of the labeling. Furthermore, it is not clear how we could ever experience the labeling, since this seems like a perceptual product, and Dennett has argued that there is no part of the brain dedicated to perceiving such products.

Also problematic is that labeling appears to be at odds with the principle. As we have just seen, he argues that there is no perceptual product. However, labeling does appear to be just such a perceptual product. It is unclear what work labeling could do if we fail to have the mechanisms necessary to experience it.

Give these considerations, it seems that the best way to preserve the most of his account is to simply reject the labeling feature. When we exclude the labeling feature, the account is no longer internally incoherent. The ignoring feature and the principle are consistent with one another. Furthermore, this allows him to maintain his Thrifty Producer Principle, which plays an important role in his attack on the Cartesian Theatre conception of mind.⁵⁸

⁵⁸ This is one of the central moves of his position in general—he has been a long-standing opponent of such a theory of mind. (See ‘Consciousness Explained’, for example)

3. REEVALUATING THE ASSUMPTIONS THAT MOTIVATE THE GRAND ILLUSION HYPOTHESIS

3.1 WHAT ASSUMPTIONS MOTIVATE THE HYPOTHESIS?

Since our attempt to give an account regarding how it is that the Ordinary Perceiver's vision seems photograph-like under normal circumstances, but allows her to experience her flaws under special circumstances, seems fruitless, it will be useful to critically review the steps that lead us to this point. It seems that if we could somehow show that our conception of predicted (non-illusory) nature of perception was such that, if our expectations were different (that is, such that we wouldn't expect to experience the flaws under normal conditions), then we would no longer have the burden of explaining why we don't experience them normally but do experience them under certain circumstances. In other words, since the Grand Illusion hypothesis is motivated by the expectations generated by our assumptions (regarding the affect that the flaws would have on our perceptual experience), by reevaluating these assumptions we may find a way of altering our expectations, such that we no longer need to offer an account as to how this illusory experience of failing to experience the flaws is possible.

Recall that Dennett assumes that the surprise exhibited by the Ordinary Perceiver when she discovers her flaws indicates that she was expecting that her experience would be different from that what she actually experienced in this case. From this, he appears to assume that, since she was surprised by her experience of the flaws (the experience of which fails to be photograph-like), then she was expecting her experience to seem photograph-like, and she had such expectations because her previous experience had seemed photograph-like to her. However, this conclusion does not follow. It seems that it is reasonable to infer that her surprise only indicates

that she was not expecting to experience what she had when she discovered a flaw. In other words, since she was surprised to experience her vision as failing to seem photograph-like, all we can infer is that she was not expecting for her experience to not seem photograph-like, but from that it does not follow that she did expect to experience to seem photograph-like.

One might wonder why I argue that Dennett implicitly assumes (or, at the very least, *inadvertently* assumes, such that it is entailed by his position) that she is expecting that her experience will be photograph-like. The following seems to be a reconstruction of how he determines her expectation. Consider the role that the surprise performs in determining whether the Ordinary Perceiver is subject to the Grand Illusion. Since she is surprised to discover her flaws, Dennett infers that she was expecting something else, namely that her experience would be photograph-like. This expectation must have been caused by a belief that her experience would be so—and this belief must have been generated by previous experience(s) which did seem photograph-like to her (at the time that the belief was formed and throughout the duration which she held the belief). He happily notes that her experience did seem photograph-like to her, and he concludes this by taking as evidence the surprise that she exhibits when she experiences her flaws for the first time. The only way that his conclusion (that her previous experience seemed photograph-like to her) follows from the surprise is with the inclusion of the missing premise that the overthrown belief is that her vision seems photograph-like, such that when she experiences her vision as not-photograph-like (during the discovery of her flaws), it is *that* belief that is overthrown. Otherwise, all he would be warranted in concluding from the surprise is that she was not expecting the not-photograph-like experience, but this belief could be generated by more than just the photograph-like experience. Whereas, the only belief

(overthrown upon discovery of the flaws) that would support his claim that her experience seemed photograph-like would be the experientially-acquired belief that her vision seems photograph-like, such that the “something else” that she was expecting was that her experience of the region of, say, her blind spot, would not be not-photograph-like (since she experiences that area during normal perception as photograph-like).

The surprise exhibited by the Ordinary Perceiver could be interpreted as indicative of both (by which I mean her expectation that her experience would be photograph-like or that it was not the case that she expected her experience to not be photograph-like would both cause surprise). However, in the case that she was not expecting that her experience would not be photograph-like, the Ordinary Perceiver cannot be said to suffer from the Grand Illusion, because we can only establish that it is not the case that her vision seems to her as though it is not everywhere photograph-like. But this is different from seeming to the Ordinary Perceiver that her vision *is* everywhere photograph-like. These beliefs are functionally equivalent (they cause the Ordinary Perceiver to behave in the same way—at least regarding the display of surprise), but their content is distinct. However, from the fact that it does not seem to the Ordinary Perceiver that her vision is not everywhere detailed, distinct, continuous, in color, etc., it does not follow that her vision does seem to her to be everywhere detailed, distinct, continuous, in color, etc. Her vision could appear many different ways and still satisfy the criterion ‘does not seem to her as though her vision is not everywhere detailed, distinct, continuous, in color, etc.’ Only if we can make the case that she has the belief_c that her vision seems everywhere detailed, distinct, continuous, in color, etc., are we warranted in inferring that it does seem to the perceiver as though her vision is everywhere photograph-like.

Consider this analogous case: our perceiver is walking through her local campus, when she gets hit with a water balloon. She is surprised by this, but why? Is it because she is expecting to not get hit with a water balloon? Presumably not--she could be surprised by this event even if she has no previous experience with water balloons (such that she is not aware that such things even exist). If surprise was only possible in cases where one was expecting the exact opposite event from the one experienced (in this case, if the perceiver could only be surprised if she was expecting to not get hit with a water balloon), then the only circumstances under which one could be surprised would be those which one is both aware of the source of the surprise (minimally, that they know such a thing exists) and expecting that their experience will be such that the source of the surprise will be absent. However, this entails that, by definition, the Ordinary Perceiver is not surprised by her experience of the flaws. This is because, due to the criterion of being unacquainted with the flaws as a prerequisite for being Ordinary Perceiver, she does not know that the source of her surprise exists, such that she could not possibly expect not to experience it.

One might protest that the advocates of the Grand Illusion hypothesis never argue that she was expecting the exact opposite of her experience and, as such, my claim that it is entailed by their position is unjustified. While it is true that no party explicitly argues as such, my conclusion is still entailed by their position. One can see this by looking at how they interpret the surprise. From the Ordinary Perceiver's surprise, they argue that she was expecting something else—which ultimately turns out to be a continuation of her previous, "photograph-like" experience. What is of interest here is how they get from her surprise to the conclusion about the "photograph-like" quality of her experience. Since she was surprised because her experience did not conform to her expectations (which are generated by a belief that is

overthrown when she experiences her flaws for the first time), it appears to be implicitly assumed that her belief was caused by her experience being "photograph-like" due to the widespread discussion regarding why it is that she fails to experience her flaws. To put the point another way, it is unclear why there would be so much debate over how we should account for her failure to experience the flaws if it is not assumed that, under normal conditions, she perceives the area which contains the flaws (but nonetheless fails to experience them). And, of course, I am not arguing that the Ordinary Perceiver analyzes her experience in this way. But this is in contrast to the belief that is requisite for warranting our attribution of the Grand Illusion to the Ordinary Perceiver. If it is not the case that her vision seems to her as though it is not everywhere detailed, distinct, continuous, in color, photograph-like, etc., this is different from seeming to her that her vision *is* everywhere detailed, distinct, continuous, in color, photograph-like, etc. The inference-based belief does not include any positive assertion about what her experience *is* like (only what it is *not* like); numerous different positive assertions about the nature of her experience are consistent with this tacit, inference-based belief. Since this belief does not have the requisite content, we cannot attribute the Grand Illusion to her. The onus is on the advocate of the Grand Illusion to give an account of how the Ordinary Perceiver can experientially acquire the belief that her vision is photograph-like that does not run into the problems already raised.

Given that these expectations are what motivate the claim that the Ordinary Perceiver is subject to an illusion, it seems plausible that by changing our assumptions, we may be able to eradicate the problem. To figure out how to change our expectations, we must carefully determine what they are, and see what assumptions led to these expectations. Once we have done this, we can challenge

those assumptions to determine whether we ought to adopt new assumptions which would alter our expectations in the desired way.

3.2 A DISTINCTION BETWEEN KINDS OF FUNCTIONAL FLAWS

It seems that the best place to start our evaluation of the assumptions that motivate the Grand Illusion hypothesis is right back at the beginning, with the functional flaws. As noted earlier, there are many different flaws, the failure to experience all of which is taken to be indicative of the Grand Illusion. However, we can make a useful distinction between two kinds of flaws based on the features of the visual system which generate them.

As noted above, part of our explanation must include attention to the predicted nature of perception. I believe that the predicted nature of perception, as understood in the Grand Illusion, rests on an invalid argument. Recall that, by the predicted nature of experience (as formulated within the Grand Illusion hypothesis), I mean the nature of perception as including the experience of the functional flaws. Given that it is assumed that perception that fails to include experience of the flaws is illusory, then presumably, if our experience included the flaws, we would not be subject to an illusion. This is derived from an argument like the following:

4. We have functional flaws, which we can experience under certain conditions.
5. We do not experience these flaws under normal conditions.
6. Since the Ordinary Perceiver has not yet experienced these flaws (and, upon doing so for the first time, exhibits surprise), we may infer that she was expecting something else to happen—thus overthrowing her belief that caused her expectation to the contrary (that her vision is detailed, distinct, continuous, in color, photograph-like, etc.). She is subject to an illusion—she does not experience her

vision as it *really is*.

However, it does not follow from the fact that we have these functional flaws which we are able to experience under certain conditions that we must notice them under other or all conditions. Furthermore, it does not follow that the predicted nature of experience which motivates our expectation that we would experience our flaws under normal conditions is based on the best possible set of assumptions that could be formulated based on the data from the relevant perceptual research.

This appears to be a problem with our expectations: because we *know* that the flaws exist, we *expect* that we should experience them. However, it is not clear that we are warranted in having such expectations. As noted above, from the fact that we have these flaws which we can experience under certain conditions, it does not follow that we ought to experience them in any other or all circumstances.

As we saw, the Grand Illusion hypothesis entails an incoherent formulation of the predicted nature of perception. Though never explicitly stated, this formulation can be directly derived from that which is taken to be illusory about the Ordinary Perceiver's visual experience. The error ultimately stems from considering the data from the two different kinds of flaws independently, then formulating expectations for each kind without considering the data of the other. When we simultaneously evaluate the two contrary accounts of predicted nature of perception, it is apparent that the formulation is incoherent.

A distinction can be made between two kinds of functional flaws. The differentiating feature between the two kinds is the aspect of the perceptual system that causes them. One kind is generated by the functional structure of the visual system. Let us call these the 'structural flaws'. The other kind is generated by the attention-dependent nature of perception. We will call these 'scope flaws', as these

flaws restrict the scope of that which we notice about the changing world before our eyes.

The structural flaws are caused by the functional structure of our visual system. By functional structure, I mean how the structure of our visual system impacts the quality of our visual information. Perhaps the best example of this kind of flaw is blind spot. The blind spot is created by the lack of photoreceptors in front of where the optic nerve connects to the back of the eye. Since there are no photoreceptors, we fail to receive visual information from the region.

Other examples of the structural flaws include the saccades, uneven distribution of rod and cones (rods primarily external to foveal region, cones restricted to foveal region) and a lessening in general of photoreceptors from the center of the visual field (where there is a high density of cones) towards the periphery (where there are no cones and very few rods).

It is important to understand the impact of the structural flaws both on her visual experience under normal conditions and also under contrived conditions which do not conform with our normal perception. The discrepancy between how the flaws impact our vision under normal circumstances and under special conditions leads the advocates of the Grand Illusion hypothesis to argue that our experience under normal circumstances is illusory. The assumption appears to be that, since we know that we can experience these flaws, we must be subject to illusion when we fail to experience them. In other words, since the functional structure of our visual system is responsible for her visual experience, then one must be subject to illusion if one's visual experience fails to conform to that which the functional structure would support. Thus, it is useful to discuss the effect each flaw has on our visual experience under both kinds of circumstances, according to the Grand Illusion hypothesis.

As noted earlier, the blind spot results from a lack of photoreceptors. Under normal conditions, we fail to notice our blind spot. However, we can easily experience our blind spot under contrived conditions, such as shutting one eye while fixing the gaze of the other on some fixed point, then attending to that region of the visual field (this usually requires some sort of image or symbol which we can observe to be altered by the blind spot—disappearing behind blind spot, etc.).

The uneven distribution of rods and cones results in differing quality of visual information from different areas of the retina. Cones are the photoreceptors that allow for color vision. They only work well in well-lit environments; they are functionally impaired in poor lighting. The high density of cones in the foveal region creates a highly detailed, richly colored visual experience of the center of our visual field. However, we have few rods in the area, making it very difficult to acquire visual information from that region in low lighting conditions.⁵⁹ Outside of the foveal region, we fail to have color vision. This is because rods, which are responsible for our ability to see under poor lighting conditions, do not support color vision.

Given the uneven distribution of, and differences in kind of visual information processed between, the two kinds of photoreceptors, the advocates of the Grand Illusion hypothesis conclude that our visual experience under normal conditions, which we do not experience as lacking color and detail in areas, must be illusory. Note that we are able to experience this kind of structural flaw under contrived circumstances. For example, we cannot detect the color, and sometimes even the shape, of an object that is held at our periphery. At best, we can detect it if it is moving—otherwise, we seem to be unable to discriminate anything about it.

Unlike the functional flaws, the scope flaws are not generated by the physical

⁵⁹ This is why, for instance, you cannot see most stars by looking directly at them. Instead, you must look slightly to the side of them.

structure of the visual system; they result from the attention-dependent nature of perception. Recent research has indicated that our visual experience is restricted to that to which we are currently attending. Two examples of scope flaws are change blindness and inattention blindness. Both of these flaws are regarding our failure to notice large-scale changes and events that are occurring immediately before us. In both kinds of cases, we are unaware that we are failing to perceive any such changes and events.

3.3 FORMULATING THE PREDICTED NATURE OF EXPERIENCE

Independently of whether the advocates of the Grand Illusion hypothesis are warranted in their assumptions regarding what is indicated by the Ordinary Perceiver's surprise, let us grant them these assumptions for the sake of argument, so that we may see what these assumptions entail regarding the predicted nature of perception. To determine what is entailed, we must first formulate the predicted nature of perception with respect to each kind of flaw. We must then compare the two formulations, such that we can then see the entire account of predicted nature of experience as entailed by the assumptions of Grand Illusion hypothesis.

The Grand Illusion hypothesis appears to assume that, given what we know of the structural flaws, any visual experience which seems everywhere detailed, distinct, continuous, in color, etc., must be illusory. Suppose this were true—what does this entail about the predicted nature of perception? Since any visual experience that is “photograph-like” is illusory, and the claim that such experience is illusory is based on our expectations of how these flaws would impact our visual experience, then it follows that the predicted nature of perception would include experience of the structural flaws. For example, since our experience is said to be illusory based on the fact that we do not generally experience a lack of visual information resulting from

the lack of photoreceptors which is responsible for our blind spot, it follows that the predicted nature of experience would include such a lack of visual information.

As such, the predicted nature of perception would be highly detailed and richly colored only in a relatively small portion of the center of our visual field, with detail greatly decreasing towards the periphery. We would experience a lack of visual information in the region of the blind spot. Furthermore, our visual experience would fail to be continuous; due to the saccades, we would experience a rapid sequence of alternating black-and-white and color “snapshots”.

Regarding the scope flaws, the illusory experience is that it seems to us as though we perceive all of the changes and events that happen before us. However, as both kinds of scope flaws demonstrate, we do not notice all of them; what we perceive is constrained by that to which we are attending. Since the guiding assumption here is that it seems to us that we notice so much more of what's going on around us than we actually do, then I suppose that the predicted nature of experience would be such that it would be apparent that we only perceive that to which we are directly attending. Perhaps the predicted nature of experience might be contained to some portion of the foveal region, such that the rest would be absent.

It is worth noting that here, as elsewhere, there is an equivocation on “seems”. One could mean that our visual experience is such that it seems to us as though we are noticing much more than we in fact are. This formulation seems strange to me—I am not sure what it would be like for my visual experience to seem to me as though I am experiencing much more of what is going on before me than I actually am. Maybe what is meant here is that the Ordinary Perceiver, who, for obvious reasons, does not notice that they fail to notice some of the events happen before them, forms the

implicit belief that they notice (all of the) changes going on in front of them.⁶⁰ We can then explain the surprise by appealing to that belief. However, this won't do either—the Grand Illusion is a *perceptual* illusion, not a false belief. The false belief, if not generated through illusory experience, is not indicative of a perceptual illusion.⁶¹

Note that all of these assumptions about the predicted nature of perception are derived directly from expectations generated by the research into these functional flaws. What is interesting is that the two kinds of flaws are never discussed together. They are either discussed in different sections of the same paper, or discussed in different papers altogether. As such, we can assume that the assumptions for each kind of flaw, regarding why failure to experience them is illusory, are independent. This is so because the case for each kind of flaw counting as evidence for the Grand Illusion hypothesis is made independently of any similar arguments for the other kind of flaw. Since the case for each kind of flaw is made independently, this entails that each kind of flaw has a distinct formulation of the predicted nature of perception based on the expectations generated by the assumptions regarding how that kind of flaw would affect our visual experience.

When we consider these assumptions together, our account of the predicted nature of perception becomes incoherent. The structural flaw formulation entails that we experience a lessening of color and detail towards the periphery, alternative sequences of black-and-white and color “snapshots”, no color outside the foveal region, etc., which requires the ability to perceive all of these areas of our visual field at once—which is exactly what the scope flaw formulation prevents by restricting our

⁶⁰ This is the kind of banal surprise that Jonathan Cohen attributes to the ordinary perceiver at the end of *The Grand Grand Illusion Illusion*.

⁶¹ See my earlier discussion of why this is so. Also, while it is interesting question regarding how it is that the Ordinary Perceiver has come to have such beliefs (if they do have them), such false beliefs are not themselves illusions (to call them such is a misnomer—illusions are perceptual).

experience to a portion of the foveal region. Alternately, the scope flaw formulation requires that our vision is restricted to the foveal region (or whatever region to which we are currently attending—it is generally, though not necessarily, the foveal region), which is prevented by the structural flaws, which entail that we *can* experience all of the structural flaws at once. Clearly, the result is two contrary formulations. Thus, we cannot maintain both kinds of flaws simultaneously.

3.4 ALTERNATIVES TO THE INCOHERENT FORMULATION OF THE PREDICTED NATURE OF EXPERIENCE

We have a few options. We could reject one of the kinds of flaws, establish that there are actually *two* (or more) distinct illusions, or reject the Grand Illusion hypothesis in the hope that some promising new direction might be found upon its deconstruction. Ultimately, the third option is the only viable one, because there is no principled reason for selecting one kind of flaw and rejecting the other, since they are both taken to be illusory. Furthermore, arguing that there are at least two distinct illusions will also fail because, ultimately, they are illusions of the same perceptual system, such that they will end up with the same predicted nature of perception, whether they are considered as one illusion, or separately.

It is worth being thorough in our inquiry regarding whether we could consider our failure to notice one of the kinds of flaws as an illusory experience, but not the other kind. Since it is only when we consider the together that the formulation of predicted nature of perception becomes incoherent, it seems reasonable to assume that, if we could successfully exclude one of the kinds of flaws, then formulation would no longer be problematic (at least with respect to internal coherence). In order to be able to differentiate the two kinds of flaws in this way, we need to find some theoretical grounds for arguing that the feature in virtue of which one formulation is

illusory is absent from the other formulation. This means that we must evaluate on what grounds we are assuming that each are illusions.

Let us begin with the structural flaws; recall that the expectation which motivates the claim that our failure to experience the flaws is illusory is generated by both our knowledge of the functional structure of the visual system as well as our ability to experience in some of these flaws under contrived circumstances. Thus, the claim that our experience under normal conditions (during which we fail to experience these flaws) is illusory stems from the assumption that, since these flaws are generated by the structure of our visual system, and the functional structure of our visual system is what processes our visual information, then it seems that this “hardware” would determine our experience.

Similarly, the intuition that our failure to notice the scope flaws (that attention restricts what we perceive) indicates that we are subject to an illusion rests on the failure of our visual experience to conform to the apparent constraints of the operations of our visual system. It is argued that it seems to us as though we notice so much more than we do about what is going on before us, and it is in this sense that our experience is said to be illusory.

Assume for the sake of argument that, if these are illusions, they are so for the reasons argued. We must compare the reasons that each are considered to be illusions in an attempt to determine whether there is any way that we could differentiate the two kinds, such that one is an illusion but the other is not. In the structural flaws formulation, we are subject to an illusion because our experience does not conform to what we would expect (given what we know about the structure of the visual system). Likewise, in the scope flaw formulation, we are subject to an illusion because our experience does not conform to our expectations (given what we know about the

attention-dependent nature of perception). Both cases share a common formulation: the perceiver is subject to an illusion because their experience does not conform to what we would expect their experience to be, given the kind of flaw under discussion. Furthermore, to deny that one is an illusion, is to argue that it cannot be an illusion on these grounds (it follows that if it is not an illusion, then it is not an illusion for these reasons, as well as any other reasons). But this would serve as a counterexample to the formulation on which the other illusion depends, such that it would also rule out the other kind as an illusion. As such, we cannot maintain that one is an illusion, but not the other.

Nor can we argue that, rather than both being parts of the Grand Illusion, they are two separate illusions. Suppose that we wish to argue that they are distinct illusions, in an attempt to avoid the problematic formulation of the predicted nature of perception when the two are considered part of the same illusion. The problem with this is that, ultimately, they are illusions of the same perceptual system, such that, if both are illusions, then the same predicted nature of perception is entailed—they problematic formulation would still arise.

It appears that our assumptions will not allow for a coherent formulation of the predicted nature of perception. As such, it appears that we ought to abandon the Grand Illusion hypothesis in search of a new set of assumptions which will allow for a viable formulation of our expectations regarding predicted nature of perception.

4. A NEW FORMULATION OF THE PREDICTED NATURE OF EXPERIENCE

4.1 CONSIDERING THE AFFECT THAT EACH KIND OF FLAW HAS ON OUR EXPERIENCE OF THE OTHER KIND OF FLAW

Suppose for the sake of argument that my deconstructive project is successful—where do we go from here? In other words, though deconstructive projects are useful in their own right, it may be fruitful to search the ruins of the old problem for insight into direction for new position. This appears to be just such a case.

Recall my argument for why we cannot coherently formulate the Grand Illusion hypothesis. I alleged that the Grand Illusion hypothesis resulted from faulty expectations regarding the predicted nature perception based on considering only isolated bits of data. When considering the data from the structural flaws, we assume that since we can experience, say, the blind spot, then our failure to do so under normal circumstances must be due to our being subject to illusion. Likewise, from the research on the scope flaws, we infer that, since the Ordinary Perceiver is surprised when she discovers that she fails to notice large, non-occluded events that are happening immediately before her, it must be the case that it seems to her as though she is receiving visual information from the entirety of her visual field (that is to say, from outside of the restricted area to which she is attending), when in fact this must be illusory, since we know she does not have any such experience of information external to the area to which she is attending through her failure to notice such events. From this, we have generated two contrary sets of expectations regarding what the predicted experience of perception must be like, such that deviation is grounds for arguing that such an experience is illusory. However, as we

saw earlier, this results in an incoherent formulation of predicted nature of perception. When we consider the two individual formulations together, we discover the two cannot be maintained simultaneously.

This failure is insightful, however; by examining what went wrong, we are able to see an alternative formulation of our expectations regarding our experience, given the functional flaws. The most important lesson here is that we must consider the data from the research on both kinds of functional flaws simultaneously in order to more accurately develop our expectations. The attention-dependent nature of perception is equally important of a consideration as the functional structure of visual system. For example, though we receive no visual information from the region which falls within the blind spot, given the attention-dependence of vision, namely we only notice that to which we are attending, which usually (though not necessarily) is that which falls within our foveal region, as well as the fact that are blind spot falls outside of that region, then we should not expect to experience it under normal conditions.⁶² Likewise, though it is true that we fail to notice large scale events that happen immediately before us, this is exactly sort of thing we should expect, given what we know about the attention-dependent nature of perception; if we are not attending to the change, we fail to notice it.

The previous incoherent formulation appears to arise due to emphasis on only a small portion of the causal sequence of perception. These two kinds of flaws are distinct with respect to their causal role in perception. In other words, there is more to consider with respect to what we experience than the means by which we process the visual information; quite literally, we do not “read it off the retinal image”. Instead,

⁶² Though it does seem reasonable to expect that we can experience it under conditions in which we are attending to that region. For example, we are able to experience our blind spot when we have closed one eye and, though we keep our gaze steady, we are attending to some image or symbol within that region.

there are other factors that determine our experience as well.⁶³

We can now begin to see a new formulation of our expectations when based on simultaneous consideration of the affects of both kinds of flaws on perceptual experience. If we note that the attention-dependent nature of perception prevents us from noticing that to which we are *not* attending, and we are (almost) always attending to the foveal region (which we *do* experience as detailed, distinct, in color, etc.), then it follows that we should expect that we would fail to notice the lack of detail, color, etc., outside the foveal region. As such, it appears as though our experience does conform to our expectations, and, thus, we are not subject to a Grand Illusion.

Once we have amended our expectations of what our experience must be like based on simultaneous consideration of both kinds of flaws, we see that our experience no longer fits the old framework for an illusion: our experience no longer fails to conform to what we would expect (given not only our knowledge about the flaws, but considerations of how they impact one another).

4.2 EXPLAINING THE SURPRISE WITHIN THE NEW ACCOUNT

However, I still owe an explanation for why it is that the Ordinary Perceiver is surprised when she experiences her flaws. First, remember that under the new formulation, visual experience conforms to our (the theorists') expectations. It does not follow from that, that the Ordinary Perceivers expectations will be the same as those which we have come to have the virtue of reevaluating the assumptions of the Grand Illusion hypothesis. She will come to form her expectations in a different way, namely through how her experience seems to her. We are now in a position to

⁶³ Note that I am *not* suggesting that there is some ultimate perceptual product—some finished object or image which is then perceived. Rather, I am arguing that there is more involved in determining what is the nature of our perception than the narrow evaluation of the flaws suggests.

understand why, though she is not subject to the Grand Illusion, her visual experience is such that we are nonetheless tempted to call it "photograph-like": perceptual experience under normal conditions is, due to the attention-dependent nature of perception, restricted to the visual information acquired through the foveal region. This region contains a high density of cones, allowing for her experience to be highly detailed and colored. Note that I am not positing any mechanism which somehow prevents her from experiencing the structural flaws; they are unhidden at all times, available to be perceived at any time to which they are attended but, due to the attention-dependent nature of perception, she doesn't notice them under normal circumstances⁶⁴.

If our newly-formed expectations regarding the predicted nature of perception are roughly correct, then her experience is of visual information acquired through the foveal region, which is everywhere detailed, distinct, continuous, and in color. Thus, anytime her vision does not seem that way to her, she is surprised, not because she was expecting the region containing the flaw to conform to a "photograph-like" experience of perception due to the fact that she previously experienced that region as such, but rather because she was not expecting it to fail to conform with her previous visual experiences (which, though not of that region, nonetheless fail to indicate that experience of the flaw regions would be any different). Even if we want to argue that, not only did she not expect that experience, but she also had expectations of what that experience would be like, all we could say is that she expects the same quality of detail and color that she normally experiences.⁶⁵

⁶⁴ During normal perception, she is attending to the foveal region, while the structural flaws which we have been discussing are located outside of the foveal region.

⁶⁵ Given the treatment of filling-in found within at least some aspects of the debate, it is worth mentioning that once our expectations have been reformulated in the proposed way, there is no longer the burden of explaining why we do not experience our flaws during normal perception. This has two potential implications for the advocate of filling-in: (a) some may abandon this sort of explanation since we no longer need to account for our failure to experience the flaws, or (b) some may accept the

4.3 CONCLUDING REMARKS

In conclusion, rather than expecting to experience the flaws under other conditions, our expectations can be modified by considering other, relevant constraints on perception. Attention, for example, is an important constraint: we perceive that to which we attend (and, surprisingly, little else). The structure of the eye is another important constraint, namely regarding the distribution of photoreceptors. How the information is interpreted (how we represent it as being) affects how it seems to us as well. Using these three examples, I will illustrate how constraints can alter our expectations.

We (at least generally) only notice that to which we are attending. This means that, when we are purposefully trying to find our blind-spot,⁶⁶ we are attending to that area of our visual field and noticing features of it in greater detail (or lack thereof) than when we are attending to other regions. However, we are normally attending to the foveal region of our vision. The blind-spot falls outside of that region, and is thus generally outside of our attention. Therefore, the blind-spot need not be altered for us to not perceive it—we may not perceive it because, like the rest of perception, we only notice that to which we are attending. Similar lines of thought have been used in theorizing about change-blindness and inattentional-blindness, and I see no reason why we cannot apply this explanation more generally.

The distribution of photoreceptors is another important constraint—the more photoreceptors, the greater the visual detail. Since there is a greater number of

new formulation, knowing that they no longer bear the strange explanatory burden of why we are able to experience our flaws under non-normal circumstance. Either way, it seems that, in accepting the new formulation, filling-in has been vindicated from what was generally taken (perhaps erroneously, since there appears to be no discussion of such a theory in its own right, but only as a means to this particular explanatory end) its previous role in the debate—accounting for why we don't experience the flaws—which lead some members of the debate to interpret filling-in as an unreasonable position.

⁶⁶ I do not mean for this to be the only case—there are other ways of discovering your blind-spot.

photoreceptors (and a great deal of those being cones) in the foveal region, which falls in the center of the visual field, then whatever falls into that region of the visual field will be experienced in color and great detail. We generally attend to that which falls within our foveal region, so everything to which we attend is generally experienced in color and great detail. Furthermore, the density of photoreceptors (especially cones) becomes less moving in the direction from the foveal region to the periphery. This lessens the detail and color of our vision. However, we do not generally notice this lessening of detail and color *because we are not attending to that region*. If we normally attend to the center of our visual field, and the center of our visual field is experienced as being in color and having good detail, we (a) do not notice that it is otherwise elsewhere in our visual field and (b) that which we are perceiving *is* experienced as detailed and in color (this is not illusory—it appears reasonable to argue that one’s experience really is that way in this region), which may lead us to infer that our vision (perhaps entirely) is that way. This inference is not subject to the illusion, because it is not the case that we are experiencing our *vision* as everywhere detailed, distinct, continuous, in color, photograph-like, etc. Rather, we experience the region of our visual field to which we normally attend (and which *is* that way) *as* being that way, and could mistakenly infer that our entire visual field is that way (since we do not generally attend to that which falls outside of it—we just turn our heads to look if we wish to see something that falls outside of the foveal region—and, due to photoreceptor distribution, which yields poor information outside of that region, we do not experience anything to disconfirm our experience until we experience the flaws). Given these modified expectations, it seems clear that we are *not* subject to the Grand Illusion.

BIBLIOGRAPHY

- Cohen, Jonathan. 'The Grand Grand Illusion Illusion.' *Journal of Consciousness Studies*. (9:5/6) pp. 141. 2002.
- Dennett, Daniel. 'How Could I Be Wrong? How Wrong Could I Be?' *Journal of Consciousness Studies*. (9:5/6) pp. 13. 2002.
- Dennett, Daniel. 'Filling In Versus Finding Out: A Ubiquitous Confusion in Cognitive Science.' <http://ase.tufts.edu/cogstud/papers/fillin.htm>, last accessed August 12, 2004.
- Dennett, Daniel. 'Seeing is Believing—or is it?' in K. Akins, ed., *Perception Vancouver Studies in Cognitive Science*, vol. 5: Oxford University Press, 1996 pp. 158-172.
- Noë, Alva. 'Is the Visual World a Grand Illusion?' *Journal of Consciousness Studies*. (9:5/6) pp. 1-12. 2002.
- Pylyshyn, Zenon. 'Mental Imagery: In Search of a Fundamental Theory.' <http://rucss.rutgers.edu/faculty/zpimagery.htm> Last accessed August 12, 2004.

SEEING BEYOND THE ILLUSION:
DISPELLING THE GRAND
ILLUSION HYPOTHESIS

Melissa A. Ebbers, M.A.

Western Michigan University, 2005

The Grand Illusion hypothesis is a new form of skepticism about the nature of our visual experience: it *seems* to us as though our vision is everywhere detailed, distinct, continuous, in color, even “photograph-like,” but it is not. This position is motivated by developments in perceptual research, which have revealed new information about the functional structure of the visual system as well as the attention-dependent nature of perception.

My project is primarily deconstructive. I argue that the Grand Illusion hypothesis rests on problematic assumptions (motivated by the results from the relevant perceptual research), which ultimately leads to an incoherent formulation of what we ought to expect our vision to be like. By challenging these assumptions, our expectations are altered in such a way that predicts that we would have exactly the visual experiences that we do. Thus, we are not subject to the Grand Illusion.

SEEING BEYOND THE ILLUSION:
DISPELLING THE GRAND
ILLUSION HYPOTHESIS

by

Melissa A. Ebbers

A Thesis
Submitted to the
Faculty of The Graduate College
in partial fulfillment of the
requirements for the
Degree of Master of Arts
Department of Philosophy

Western Michigan University
Kalamazoo, Michigan
April 2005

Copyright by
Melissa A. Ebbers
2005

ACKNOWLEDGMENTS

First and foremost, I would like to express my gratitude to the members of my thesis committee: Dr. John Dilworth, Dr. Timothy McGrew, Dr. Michael Pritchard, and Dr. Marc Alspector-Kelly. Their involvement in this process and genuine enthusiasm for the project contributed greatly to my success.

Secondly, I would like to thank my wonderful friends—including my classmates—and family (especially my mother and father) for their love and support during this process. Thank you for all of the fun and interesting discussions, as well as your kind words of encouragement.

Lastly, I would like to express great appreciation for the hard work and dedication of my advisor, Dr. John Dilworth—he certainly deserves special recognition for his contribution to this effort. His guidance and support, which included thorough comments, frequent meetings, advice, encouragement and helping me to keep a sense of humor about the project, has made this a highly rewarding experience. I am truly fortunate to have such a wonderful mentor.

Melissa A. Ebbers

TABLE OF CONTENTS

ACKNOWLEDGMENTS	ii
CHAPTER	
I. AN OVERVIEW OF THE GRAND ILLUSION HYPOTHESIS	1
1.1 Introduction	1
1.2 The Grand Illusion Hypothesis: A New Skeptical Problem?	5
1.3 How Do We Determine Whether the Ordinary Perceiver is Subject to the Grand Illusion?	5
1.4 The Role of the Ordinary Perceiver.....	11
1.5 The Ordinary Perceiver’s Belief.....	12
1.6 What Does (and Does Not) Surprise the Ordinary Perceiver	12
II. THE PREDOMINATE METHOD WITHIN THE DEBATE FOR DETERMINING WHETHER THE ORDINARY PERCEIVER IS SUBJECT TO THE GRAND ILLUSION	16
2.1 Overview	16
2.2 Distinction Regarding Methods of Belief-Acquisition.....	18
2.3 The Method of Acquiring Experience-Based Beliefs.....	21
2.4 Other Candidates for the Belief.....	24
2.5 Further Consideration of Inference-Based Beliefs	26
2.6 Why Dennett Cannot Argue That the Ordinary Perceiver is Subject to the Grand Illusion.....	31
III. REEVALUATING THE ASSUMPTIONS THAT MOTIVATE THE GRAND ILLUSION HYPOTHESIS	38
3.1 What Assumptions Motivate the Hypothesis?	38

List of Figures—continued

3.2 A Distinction Between Kinds of Functional Flaws	43
3.3 Formulating the Predicted Nature of Experience	47
3.4 Alternatives to the Incoherent Formulation of the Predicted Nature of Experience	50
IV.A NEW FORMULATION OF THE PREDICTED NATURE OF EXPERIENCE	53
4.1 Considering the Affect that Each Kind of Flaw Has on Our Experience of the Other Kind of Flaw	53
4.2 Explaining the Surprise Within the New Account	55
4.3 Concluding Remarks	57
BIBLIOGRAPHY	59

1. AN OVERVIEW OF THE GRAND ILLUSION HYPOTHESIS

1.1 INTRODUCTION

The Grand Illusion is an illusion of which the layperson with normal vision (let us call her the Ordinary Perceiver) purportedly suffers. As characterized by Alva Noë, the Grand Illusion is a new kind of skepticism regarding how our visual experience seems to us. The advocates of this illusion argue that it seems to the Ordinary Perceiver that her vision is everywhere detailed, distinct, continuous, in color, photograph-like, etc.,¹ when her vision is, in at least some regions of her visual field, in fact poor, not detailed and even colorless!² As a means of determining whether such a disparity exists, discussion has focused around how to interpret the surprise—and lack of surprise—exhibited³ by the Ordinary Perceiver upon experiencing her functional flaws (blind-spot, change-blindness, inattentional-blindness, saccades, etc.) as well as during normal perception.⁴ This surprise is taken to be an indication that a belief is overthrown—that the Ordinary Perceiver was expecting something else—when she discovers her functional flaws. In what follows,

¹ I will also refer to this list of attributes as “photograph-like” (as an abbreviation of the list) This list of attributes is intended to be similar to the “snap-shot” conception addressed by Noë (*Is the Visual World a Grand Illusion?*, p.2)

² This assumption, that the predicted nature of vision is not everywhere detailed, distinct, continuous, in color, photograph-like, etc., is based on a further assumption that experiencing the functional flaws is the what we ought to expect that our visual experience would be like. I challenge this assumption in Chapters 3 and 4.

³ As discussed in §1.3, I am only arguing against the predominant method (in the literature on the debate) of determining whether the Ordinary Perceiver is subject to the Grand Illusion—through interpretation of the significance of the surprise, and lack thereof, exhibited by her under difference circumstances that are relevant to the debate.

⁴ Noë argues that it is also important to note that the Ordinary Perceiver is *not* surprised by some of the adjustments (putting on her glasses, relocating, squinting, looking around, etc.) she must make to in order to perceive the desired information during normal perception. I address this in greater detail below.

I argue that the Grand Illusion hypothesis is motivated by expectations of what our visual experience must be like which are formed without consideration of sufficient relevant factors. Since these expectations are not met during our normal⁵ perceptual experience, as taken to be particularly evident given the surprise exhibited by the Ordinary Perceiver when she experiences any of her functional flaws, we assume that we are subject to an illusion. However, the surprise alone is insufficient to settle the matter since more than one belief could cause the surprise. To illustrate this point, let us discuss what motivates us to call the Ordinary Perceiver's experience illusory, namely expectations derived from recent research of the flaws.

Since the Grand Illusion hypothesis rests on expectations of what our perception must be like which are developed without considering all of the relevant factors, it is useful to attempt to formulate the predicted⁶ nature of perception based on what the Grand Illusion hypothesis takes to be illusory about the Ordinary Perceiver's perceptual experience as a means of drawing attention to the problematic assumptions which generate these expectations. My account, then, is two-fold: the first half is dedicated to what can be said regarding how the Ordinary Perceiver's visual experience seems to her, and the second discusses the implications of arguing that the Ordinary Perceiver's experience, as formulated by the Grand Illusion hypothesis, is illusory, concluding that the Grand Illusion hypothesis is motivated by unwarranted expectations (based on considering too little of the relevant data discovered through perceptual research), and arguing that once we generate our expectations by considering the affect of each of the kinds of flaw on our experience of the other, we can then see that we are not subject to an illusion—in fact, what we

⁵ By 'normal,' I mean how our experience seems to us during most, if not all, of our experiences—all but those occasions under which we perceive our structural flaws (which we do by means of non-standard usage of our perception)

⁶ See discussion of this term in §1.3.

experience is just what we would expect to experience, once we consider enough of the relevant (available) facts.

My inquiry during the first half of my project primarily addresses how the Ordinary Perceiver's vision seems to be, and incorporates other aspects of the debate only insofar as they aid in this project. I introduce the skeptical problem in Chapter 1 by addressing the functional flaws (§1.3) and what does (and does not) surprise the Ordinary Perceiver (§1.6). Through my analysis of how the Ordinary Perceiver's vision seems to her in Chapters 1 and 2, I discuss at length the role of the Ordinary Perceiver (§1.4), specifying what belief the Ordinary Perceiver must have in order to be subject to the Grand Illusion (§1.5-2.3) and whether there are any other candidates for the belief that causes the surprise (§2.4-2.5). I conclude the first half, by arguing that the case cannot be made for the alleged disparity between how the Ordinary Perceiver's vision seems to her (everywhere detailed, distinct, continuous, in color, photograph-like, etc.) and the predicted nature of her vision (discontinuous, colorless and without detail in some regions, etc.), and, as such, a new account is needed to explain the Ordinary Perceiver's reaction to various conditions. Since we cannot develop an account to explain this disparity of experience, in Chapter 3 I suggest that a new approach is adopted by the debate—re-examining our assumptions in an attempt to discover whether we may alter our expectations in such a way as to eliminate the Grand Illusion hypothesis in favor of a new explanation of our perceptual experience.

The second half of my project, then, is dedicated to what the Grand Illusion hypothesis implies about the predicted nature of perception (§3.1), which, if I am right, renders the set of expectations (generated by the assumptions made about the results of the research on the functional flaws) incoherent, illuminating an internal

problem with the formulation of the hypothesis (§3.3). In order to draw out the incoherence, I make a distinction between two kinds of functional flaws based on how they are generated (§3.2). Since it is claimed that the surprise exhibited by the Ordinary Perceiver when she experiences each kind of flaw (which are distinct kinds since they are generated by separate aspects of the perceptual system), then, since they are argued separately to be indicative of the Grand Illusion hypothesis, they admit of a distinct formulation of both how they seem to the Ordinary Perceiver, as well as of the predicted nature of perception. In other words, once the distinction is made between the two kinds of flaws, we can evaluate what the claim that the Ordinary Perceiver is having an illusory experience implies for the formulation of the predicted nature of perception regarding each kind of flaw. The two formulations of the predicted nature of perception are compared in an attempt to integrate them into one unified formulation of the predicted nature of perception. However, I argue that, once we attend to both kinds of functional flaws, the formulation of the predicted nature of perception becomes incoherent (§3.3).

Though my account is primarily deconstructive, I conclude by arguing that direction can be found for a new formulation of our expectations and the predicted nature of perception (§4.1). Once we consider both kinds of flaws (and the effect that they will have on our overall perceptual experience), we are able to formulate one account of the predicted nature of perception (rather than two separate formulations—one for each kind of flaw, and is based solely on the expectations that we would generate by virtue of that which we know about that flaw independently of how any other feature(s) of perception might effect it—which cannot be integrated into a single account). The expectations generated by the new formulation of the predicted nature of perception are such that we need no longer entertain the possibility that we are

subject to an illusion—this account entails that we should not expect to experience our flaws during normal perception (§4.3).

1.2 THE GRAND ILLUSION: A NEW SKEPTICAL PROBLEM?

Noë argues that the Grand Illusion hypothesis is a new skeptical problem: skepticism about our perceptual experience.⁷ The worry is whether our perceptual experience is different than it seems to us. The alleged disparity is that our vision seems to us to be everywhere detailed, distinct, continuous, in color, photograph-like, etc., but this experience does not accord with the predicted nature of vision, which is not everywhere detailed, distinct, continuous, in color, photograph-like, etc., due to the impairing effects of the functional flaws on our visual system. This skepticism is motivated by the surprise exhibited by the Ordinary Perceiver when she discovers one of her functional flaws. Since she is surprised, we infer that a belief was overturned—a belief, the contents of which lead her to expect that her vision would seem to her as it always had: detailed, distinct, continuous, in color, photograph-like, etc. In other words, the surprise at discovering one of her functional flaws indicates that her belief about the photograph-like nature of her vision was overturned at that time. Thus, we worry whether we ought to be skeptical about our visual experiences, given the alleged disparity between the predicted nature of our visual experiences and how they seem to us.

1.3 HOW DO WE DETERMINE WHETHER THE ORDINARY PERCEIVER IS SUBJECT TO THE GRAND ILLUSION?

Throughout the present critique of the Grand Illusion hypothesis, I restrict the scope of my evaluation to the predominate method within the debate for determining

⁷ Noë, *Is the Visual World a Grand Illusion?*

whether the Ordinary Perceiver is subject to the Grand Illusion.⁸ As such, I take issue with using the surprise exhibited by the Ordinary Perceiver when she first discovers her flaws as indication that she is subject to the Grand Illusion. In order to analyze the Grand Illusion hypothesis in the proposed fashion, we must first introduce some terms we will be using to discuss the debate. Those terms are used to must lay out the problem. Once we have clearly formulated the problem, we will be in a position to critically review the it.

Since it is my contention that the Grand Illusion hypothesis is motivated by a disparity between our expectations regarding what we would expect (based on the recent results from perceptual research), I will define my terms to reflect this. I propose the term ‘predicted nature of experience’ to refer to what we would expect our experiences to be like (relative to some set of data from which we form our expectations). Whatever assumptions regarding what we expect that our experience would be like (that we make based on the data that we are considering) are collectively considered to be the predicted nature of experience.

In contrast, an illusion (as the term is, arguably, used within the debate) is any experience which deviates from the respective predicted nature of experience. Though this is a non-standard usage of the term, something like this appears to be the implicit use of illusion within this debate. For example, given the data regarding the functional structure of the visual system (blind spot, photoreceptor distribution, etc.), we expect that our experience would be missing visual information in the region of the blind spot and would noticeably lessen in detail and color towards the periphery; these expectations comprise the predicted nature of experience. By virtue of the apparent failure of our normal experience to conform to our expectations, such

⁸ Thus, I am not speaking to other ways of arguing for the problem. I am only concerned with the general way in which the hypothesis is taken to be established.

experience is taken to be illusory.

The **functional flaws**⁹ of our visual system are aspects of the functional structure of our eyes that lead us to expect that our vision would not be everywhere detailed, distinct, continuous, in color, photograph-like, etc. The assumption appears to be that, given the existence of these flaws, we would expect that our vision is experienced as only having rich color and detail in the foveal¹⁰ region, lessening in color and detail towards the periphery, rapid sequences of snapshots, etc. Despite the presence of these flaws as part of the functional structure of our eyes, we do not notice them during normal perception¹¹ and are surprised to discover them.¹²

I am only addressing some of the flaws commonly found in the literature on the Grand Illusion that are relevant to my discussion of the disparity between how our experiences seem to us and the predicted nature of perception, such that only those flaws which can be experienced are to be taken into account of this debate. Saccades (a consequence of the function of photoreceptors), for instance, are discussed in the literature, but since I am interested in explaining the surprise exhibited by the Ordinary Perceiver when she experiences her flaws (and I am unaware of circumstances under which she could experience her saccades), I will not be discussing them.

One of the most popular of these functional flaws, the blind-spot, is the source of much excitement (and the cause of much speculation) for the visual theorist. The

⁹ I owe this term to discussions with John Dilworth. It is important to note that I am using the term “flaw” to connote that, in the debate, it is implied that these “flaws” are that which prevents us from having (the non-illusory) experience of vision as everywhere detailed, distinct, continuous, in color, photograph-like, etc.

¹⁰ The fovea of the eye is the region which we generally use during normal vision (see footnote 8). It contains a high number of cones (photo-receptors that enable us to have color vision), which allows us strong visual discriminatory power (good visual acuity, color, etc).

¹¹ By ‘normal perception,’ I mean one’s vision when one is not attempting to experience one of the flaws—this is the kind of vision we experience most, if not all, of the time.

¹² Under non-normal perceptual circumstances, such as experimental conditions.

blind-spot, which is caused by a region that contains no photoreceptors where the optic nerve connects to the eye and is located just 6 degrees from the center of the visual field (covering 16 degrees of the visual field), can be observed under certain circumstances. These circumstances differ from normal perception: during normal perception, we attend to that which is straight ahead of where our eyes are directed, while we must attend to that which is outside of the center of our vision in order to experience our blind-spot. Whatever objects or events that fall within our blind-spot are experienced (when we are *attending* to that region of our vision) as occluded from view. One of the reasons that the blind-spot is so interesting is that, though it is not far from the center of our visual field, we fail to experience it during normal perception.

Photoreceptor density and distribution is another functional flaw. We are virtually colorblind outside of the foveal region, and the level of perceptual detail diminishes from borders of the foveal region to the periphery. We do not notice either of these flaws during normal perception (as noted earlier, our experience seems to be highly detailed and richly colored during normal perception).

Change-blindness and inattentional-blindness are two related functional flaws. Both involve the failure to experience that to which we are not attending. In an experiment on inattentional-blindness, 'Gorillas in our Mist' (Simons, D. and Chabris, C., *Perception*, 1999, **28**), subjects were asked to watch a video recording of a basketball game, attending to the number of passes made by one team. During the game, a person in a gorilla costume walks onto the court, pauses to beat his chest, then exits the court. Though many of the participants correctly counted the number of passes, fewer than half noticed the gorilla. The participants were shocked when they watched the video again, observing the gorilla which they had previously not noticed.

Surprising results are also obtained through research on change-blindness.

Subjects fail to notice changes in stimuli that occur immediately in front of them. One such example is stimuli with “mud-splashes”: the subjects observe a picture, which changes every time a “mud-splash” (a shape that occludes part of the picture briefly) flashes.¹³ Subjects are surprised to learn that they failed to notice changes that occurred in the picture as they were looking at it. This sense of surprise—common to the discovery of change-blindness, inattentional-blindness and the blind-spot (as well as others)—acts as a basic step in arguing for the Grand Illusion.

Now that we have discussed the terms, we are in a position to lay out the problem. Recall that I argue that the Grand Illusion hypothesis is motivated by a deviations from the predicted nature of experience (as formulated based on the results from the research on the functional flaws). Thus, in order to determine whether she is subject to the Grand Illusion, we must answer the following questions:

- (1) Is there a discrepancy between the predicted nature of our experience and how our experience seems to us?
- (2) What is the predicted nature of our perceptual experience?
- (3) How does our perceptual experience *seem* to us?

The main question of the debate is (1). If the answer to (1) is no, then the Grand Illusion hypothesis (at least insofar as it accords with our formulation) is false; otherwise, further analysis is warranted. The answer to (1) will become more apparent in light of the answers to (2) and (3). Question (2) asks about the predicted nature of experience. According to the Grand Illusion, our vision, due to its functional flaws, is not everywhere photograph-like. In answering (2), we must address what facts about us (and the world) determine the predicted nature of vision, such that we are justified in asserting that our expectations regarding the predicted nature of perception are such

¹³ J. Kevin O'Regan offers an excellent example of change-blindness experiments that you can try for yourself at his website (<http://nivea.psycho.univ-paris5.fr/>).

that, when contrasted against the Ordinary Perceiver's experience (allegedly, that is it photograph-like), we are warranted in saying that she is subject to an illusion. Question (3) asks how experience really seems to the Ordinary Perceiver. The Grand Illusion hypothesis alleges that vision seems to the Ordinary Perceiver as though it is everywhere photograph-like. For present purposes, we must inquire how vision seems to the Ordinary Perceiver—specifically, whether it seems to her as though her vision is everywhere photograph-like. We may discover that there is no disparity between how our experience seems to us and the predicted nature of vision, in which case (1) is false. The only circumstances under which perceivers are subject to the Grand Illusion are if the answers to (2) and (3) pick out unique experiences/events/states/phenomena.¹⁴ If this is the case (if (1) is true), then we are justified in our skepticism at least to the extent that further analysis is reasonable. However, one further condition must be satisfied. Not only must there be a discrepancy between the answers to (2) and (3), but it must be the discrepancy specified by the Grand Illusion: it must seem to perceivers as though their vision is everywhere photograph-like, even though this is not the case.

Since we are concerned as to whether the Ordinary Perceiver is subject to the Grand Illusion, we need only inquire as to whether her vision seems to her to be everywhere photograph-like. If her vision does seem to her to be this way, then we are at least warranted in further investigation regarding the Grand Illusion.¹⁵ If it does not seem this way to her, then we have answered our question (at least regarding the Grand Illusion as formulated in this paper—which captures the skepticism discussed by Noë)—the Ordinary Perceiver does not suffer from the Grand Illusion.

¹⁴ I am using neutral language here to be inclusive to different theories of mind.

¹⁵ We would still need to determine whether the predicted nature of nature of experience is as it is claimed to be (in the formulation of the Grand Illusion).

1.4 THE ROLE OF THE ORDINARY PERCEIVER

One might wonder what are the conditions under which an individual may be counted as an Ordinary Perceiver, and why she is important to the debate. Surely her perceptual capacity is not so different from our own—what warrants preferential treatment for *her* beliefs? The Ordinary Perceiver is an individual who is not acquainted with visual theory. She is not aware that the structure of her visual system is such that it includes the functional flaws (mentioned above), and she is not part of the community that theorizes¹⁶ about perception. Dennett argues¹⁷ that she possesses “pre-theoretical” views about the nature of perception, based solely on how her experience seems to her. Let us grant for the sake of argument that those who argue as such are right—that she does have such commitments. These pre-theoretical views would be valuable, presumably, because they would not be biased by knowledge of visual theory, thus serving as an unbiased means for determining how experience “really” seems. Her role, then, is to shed light on whether there is a disparity (of the sort relevant to the Grand Illusion) between how our visual perception seems and the predicted nature of vision by informing us on how her vision seems to her. Thus, we need access to her pre-theoretical views to discover how vision really seems (without the influence of theoretical bias). This is, of course, a questionable assumption; it entails that we take her perception to be entirely void of theoretical bias, including bias from *folk theorizing*.¹⁸ As Dennett notes,¹⁹ there is strong evidence in support of

¹⁶ By theorizing (here), I mean the kind that the visual theorist is concerned with—she may be part of a community that engages in folk theorizing about perception.

¹⁷ *Surprise, Surprise*

¹⁸ I have benefited from discussions with John Dilworth on this matter. In my thesis, I discuss, at greater length, the possibility of an Ordinary Perceiver. I conclude that such an unbiased perceiver is not possible—at least *some* theoretical commitments are necessary in order to “know how to” maneuver oneself in order to achieve desired ends based on the information obtained through one’s perception (the fundamental thesis of the enactive approach).

¹⁹ Dennett, *How Could I Be Wrong? How Wrong Could I Be?*

the conclusion that the layperson *does* engage in folk theorizing. If she did not, then we would expect for her to not have *any* views about the nature of her own perception. But she does seem to have such views; what else could account for her confidence in her own phenomenology and her shock at the mere suggestion that it could be otherwise?

Nonetheless, for the sake of argument, let us grant that there exist Ordinary Perceivers of the variety that the debate favors—completely without commitments to the theorizing with which the Visual Theorist concerns herself²⁰—and that a perceiver without *these* commitments is sufficient for the proposed role of the Ordinary Perceiver in resolving the matter. Thus, we assume that the impact of the folk theory is insufficient to bias the experience of the Ordinary Perceiver in such a way as to destroy her function in determining whether she suffers from the Grand Illusion.

1.5 THE ORDINARY PERCEIVER'S BELIEF

We must figure out why she is surprised to discover her functional flaws—what belief was overthrown, what expectation unmet? Since the hypothesis is that it seems to the Ordinary Perceiver that her vision is everywhere photograph-like, her belief must also have the content that her vision seems “photograph-like” (and her behavior must be governed by expectations caused by this belief) in order for us to be warranted in claiming that she is subject to the Grand Illusion.²¹ This must be the belief that is overturned in order for her to have been subject to the Grand Illusion.

1.6 WHAT DOES (AND DOES NOT) SURPRISE THE ORDINARY PERCEIVER

²⁰ That is to say, at least without theoretical commitments to the functional flaws.

²¹ Due to the length of this paper, I omit my discussion of other potential explanations of why the Ordinary Perceiver is surprised by experiencing her functional flaws, such as other candidates for the belief (other than the belief argued for by Cohen).

Surprise is taken to be the indication by which we can determine how the Ordinary Perceiver's experience seems to her: it serves as empirical evidence that her experience does not accord with her expectations (which are generated by her beliefs). Since we must discover how her experience seems to her in order to determine whether she is subject to a Grand Illusion, it is crucial to address what does (and does not) surprise her.

This is particularly of interest to us, since it seems to be the prevalent method, within the debate, of attempting to establish that the Ordinary Perceiver is subject to the Grand Illusion. If we are unable to establish the Grand Illusion in this way, then those who argue in favor of it must find a new way of grounding their position. The interesting question then becomes one of how one could otherwise establish the disparity between how the Ordinary Perceiver's experience seems and how it actually is.²² Any new method of doing so must not allow any theoretical bias to enter into her interpretation of her experience—this rules out options such as asking the Ordinary Perceiver about her experience, because the language used in the questioning might bias her interpretation of how her experience seems to her.

As noted, Noë states that Dennett is one of the main proponents of the Grand Illusion and an advocate of the position that the Ordinary Perceiver's surprise is indicative of the disparity between the predicted nature of experience and how it seems to her.²³ Indeed, Dennett does appear to endorse such a view.²⁴ He argues that the surprise exhibited by the Ordinary Perceiver indicates that a belief has been overturned—that she was expecting something else. Since her surprise was caused by discovering her functional flaws, Dennett argues that she must have had a belief that

²² By which I mean the predicted nature of experience, which is assumed to be real nature of vision.

²³ Noë, *Is the Visual World a Grand Illusion?*

²⁴ Most notably in the following articles: *Seeing is Believing*; *Surprise, Surprise*; and *Filling In versus Finding Out*

led her to expect something else would occur: namely, the belief that her vision is everywhere detailed, distinct, continuous, in color, photograph-like, etc., which led her to expect that her vision would seem that way in this case, as it had seemed to her to be in all her past experiences.

In ‘Surprise, Surprise,’ a reply to Noë and O’Regan regarding the Grand Illusion, Dennett notes that his position disagrees with that of Noë and O’Regan with respect to the beliefs held by the Ordinary Perceiver. Noë and O’Regan argue that:

“But is it true that normal perceivers think of their visual fields this way [as in sharp detail and uniform focus from the center out to the periphery]? Do normal perceivers really make this error? We think not. . . . normal perceivers do not have ideological commitments concerning the resolution of the visual field. Rather, they take the world to be solid, dense, detailed and present and they take themselves to be embedded in and thus to have access to the world.”²⁵

Dennett argues to the contrary, stating that the Ordinary Perceiver *does* hold a belief that her visual experience is everywhere photograph-like. He states that such a belief can be exposed by observing the surprise exhibited by an Ordinary Perceiver during an experiment in which she discovers, say, her blind spot. He writes:

“Surprise is a wonderful dependent variable, and should be used more often in experiments; it is easy to measure and is a telling betrayal of the subject’s having expected something else. These expectations are, indeed, an overshooting of the proper expectations of a normally embedded perceiver-agent; people shouldn’t have these expectations, but they do. People are shocked, incredulous, dismayed; they often laugh and shriek when I demonstrate the effects to them for the first time. These behavioral responses are themselves data in good standing, and in

²⁵ Noë, *Is the Visual World a Grand Illusion?*

need of an explanation. They are also, of course, highly reliable signs of their “ideological commitments”—the very commitments that elsewhere in their essay the authors correctly cite as culprits that help explain resistance to their view... Surprise is only possible when it upsets belief.”²⁶

Dennett, as noted above, argues that the surprise exhibited by the Ordinary Perceiver indicates that her experience leads her to expect something other than the experience generated by the flaws. Thus, she is subject to the Grand Illusion.

Noë notes that, though the surprise discussed by Dennett must be explained, so must *lack of surprise*.²⁷ He notes that the Ordinary Perceiver is not surprised when she must put on her glasses, relocate, look around, squint, etc., to get better visual information. It seems that she would be surprised at the need for these adjustments if she believes that her vision is everywhere detailed, distinct, continuous, in color, photograph-like, etc. It is not clear why such adjustments would coincide with her expectations (thus preventing the surprise) if her experience seemed to her as though she had all the detail at once, like a photograph. He writes:

“Surprise requires explanation, but so does lack of surprise. Notice that we are not surprised or in any way taken aback by our need to move eyes and head to get better glimpses of what is around us. We peer, squint, lean forward, adjust lighting, put on glasses, and we do so automatically. The fact that we are not surprised by our lack of immediate possession of detailed information about the environment shows that we don’t take ourselves to have all of that information in consciousness at once. If we were committed to the snapshot conception, wouldn’t we be surprised by the need continuously to redirect our attention to

²⁶ Dennett, *Surprise, Surprise*

²⁷ Noë, *Is the Visual World a Grand Illusion?*

the environment to inform ourselves about what is there?"²⁸

This implies that the Ordinary Perceiver is not committed to the belief that she experiences "photograph-like" vision.

Taking stock, then, we can see that the Ordinary Perceiver *is* surprised to discover her flaws and is *not* surprised by her need to move around, put on her glasses, look around, etc., to obtain the desired visual information. This means that a satisfactory account must include an explanation for both the surprise and lack thereof. We must still determine the methods needed to develop a satisfactory explanation of this behavior.

2. THE PREDOMINATE METHOD WITHIN THE DEBATE FOR DETERMINING WHETHER THE ORDINARY PERCEIVER IS SUBJECT TO THE GRAND ILLUSION

2.1 OVERVIEW

Let us start at the beginning and work through the steps towards the conclusion that the Ordinary Perceiver is subject to the Grand Illusion (in accordance with the approach most commonly utilized within the debate). We notice that the Ordinary Perceiver is surprised to discover her flaw(s). We take her surprise to indicate that a belief has been overthrown, and we wonder what the content of that overthrown belief must have been in order for her experience of the flaw to have been surprising to her. Surprise only occurs when our experience fails to conform to our expectations. Our expectations are generated by our beliefs, which are, in turn, caused by our experiences. As such, surprise is taken to be our outward display of the overthrowing of the belief that generated the expectations which were unmet in the

²⁸ Noë, *Is the Visual World a Grand Illusion?* pg. 7

surprising situation. Her surprise, then, was caused by the failure of her experience in this case to conform to her expectations, resulting in the belief responsible for the faulty expectation being overthrown.

For example, suppose that she experiences her blind spot. In this case, her experience is such that a portion of her visual field fails to have incoming visual information. She exhibits surprise during this experience. However, from this point forward in our evaluation of the significance of the surprise, there are two ways of proceeding. At this point, I will follow the manner used in the Grand Illusion hypothesis, to see how it turns out. Later, I will return to this point²⁹ in an attempt to formulate a new interpretation of the significance of the surprise.

In order for the surprise to be generated by beliefs which she acquired through past experience (such that her belief that her vision seems "photograph-like" was generated by her experience being "photograph-like"), her past visual experiences must have generally seemed to be everywhere detailed, distinct, continuous, in color, etc. Those who argue for the Grand Illusion hypothesis must give an account of how it is that she fails to experience her flaws under normal conditions, such that she is able to have the requisite "photograph-like" experience which could generate the purported belief that her vision is "photograph-like."

However, her surprise alone is still insufficient to warrant attribution of the Grand Illusion to the Ordinary Perceiver. While it can be argued that the surprise exhibited by the Ordinary Perceiver when made aware of their functional flaws for the first time is indicative of a belief being overthrown, the debate needs a means of explaining and resolving the disagreement between those who argue that this is sufficient evidence that the Ordinary Perceiver is subject to the Grand Illusion and

²⁹ Chapter 4

those who argue to the contrary. The crucial distinction, I would suggest, lies in the method of belief acquisition: how did she come to have the belief in question? To this end, I propose a distinction between two ways that the relevant belief could have been acquired: experience-based beliefs versus inference-based beliefs.

2.2 DISTINCTION REGARDING METHODS OF BELIEF-ACQUISITION

Experience-based beliefs (call them ‘belief_e’) are acquired through perceptual experience. Inference-based beliefs are acquired through inferences made from other beliefs. In other words, I acquire an inference-based belief, (call these ‘belief_i’) as a result of inferring the content of that belief_i from another belief (which may itself either be a belief_i or a belief_e). An example will be helpful to illustrate the distinction. Consider the following incident which happened to me recently while shopping for a new shirt. I was studying a particular shirt, commenting to my companion that I liked the shirt very much except for the pattern. My companion then drew my attention to another shirt across the table, and asked whether I liked that shirt better. This shirt had a pattern that I liked much more. I decided that it was better than the other one, since, other than the more preferable pattern, it seemed to be qualitatively identical in all other respects. It was only after I had made that statement aloud regarding my judgment that the two shirts were otherwise alike that my companion to my attention to an interesting detail that I had missed: the shirt was not, as I had thought, similar in all other respects—the collar was entirely different (the first had a button-less v-neck collar, while the second was a polo shirt with three buttons).

My belief that this shirt had the same kind of collar as the first was inferentially acquired. Was I subject to illusion in this case? Though I judged the two shirts to have similar collars, it does not follow that I had illusory experience. I made my judgment about the second shirt without having any experience of the collar at

all—it was not the case that I had an *illusory* experience of the second shirt as having the same collar as the first (such that I would have acquired a belief_e that the collars were the same style); instead, that the collar seemed to be the same was a judgment that was not based on perceptual experience, illusory or otherwise.

The proposed distinction does important work for us regarding what is meant by ‘seems to be the case to the Ordinary Perceiver.’ Discussion of how something *seems* to an individual appears (at least in this debate) to refer to the *experience(s)* that are being had by the individual. The distinction discerns between cases where it *seems* to me (experientially) that the second shirt has the same kind of collar as the first, and cases where I infer that the second collar is the same as the first (perhaps I implicitly assumed that all of the shirts on that table must have the same kind of collar, and I failed to have any experiences of the shirt that suggested that it was not so). It is only in the experiential cases that we have the “seeming” that is required for the illusion. The beliefs_i, in contrast, only tell us that it is not the case that it seems to the Ordinary Perceiver that her vision is not everywhere detailed, distinct, continuous, in color, photograph-like, etc. This is so because she is not receiving the requisite *experiential* information to form a belief regarding a positive statement about how her experience seems to her; instead, the information that she *is* receiving (primarily from the foveal region) seems photograph-like and, given no information that it is otherwise, she *infers* that the rest of her visual field is this way as well.

It is important to note the distinction between what is meant by “seems” when it is used to describe the content of the two kinds of beliefs. Regarding the content of the experientially-acquired belief, by “seems”, we mean that, to the perceiver, the content of the perceptual experience was sufficiently robust to cause the belief. By “sufficiently robust,” I mean that the nature of the perceiver's experience is sufficient

to cause the belief (rather than relying on the kind of inference that led to my belief that the shirts has the same kind of collar). In my earlier example, the experientially-acquired belief that the second shirt had the same kind of collar as the first could only have been generated by having an (illusory) experience of the shirt as having that style of collar. In contrast, “seems”, as used to describe the content of an inferentially-acquired belief, means that the perceiver judged it to be so without having actually experienced it as such. For example, though I might have been surprised to learn that the second shirt had a different style of collar than the first, and I might have exclaimed the collars seemed to be the same to me, my belief that the second collar was the same style was not acquired through a perceptual experience of the second shirt as having a v-neck collar. Therefore, it seems incorrect to say that I was subject to an illusion—rather, I acquired the belief through the fact that the part of the shirt to which I was attending was similar in style to the first, and that I was not perceiving any information to the contrary; in fact, I had not glanced at the collar at all. The fact that I had not, at any point during which my belief about the similarity in style between the two shirts was formed, looked at the collar makes it very difficult for one to argue that I was subject to an illusion—I cannot be said to have experienced the collar of the shirt as a button-less v-neck (which would have been the illusory experience) if I did not experience the collar at all.

Thus, we can see the importance of the belief being acquired experientially. This difference between the content of these beliefs (though *functionally* equivalent)³⁰ matters because, in order to attribute to her the Grand Illusion, it must *seem* to the Ordinary Perceiver that her vision is everywhere photograph-like.³¹ This “seeming”

³⁰ By functionally equivalent, I mean that both beliefs would cause the Ordinary Perceiver to exhibit surprise when she discovers her flaws.

³¹ I suppose that it could be argued that inference-based beliefs could be sufficient to motivate the Grand Illusion hypothesis, but this seems to be a non-standard use of the term ‘illusion’. Illusions are understood to be experiential and, though the Ordinary Perceiver would exhibit the surprise if the

can only be established in the case of beliefs_e. Thus, we must determine whether this belief, if held by the Ordinary Perceiver, could have been experientially-acquired such that we are warranted in attributing the Grand Illusion to her.

2.3 THE METHOD OF ACQUIRING EXPERIENCE-BASED BELIEFS

How could she acquire the requisite belief experientially? Since she could only do so through experiences that seemed everywhere photograph-like, we must determine whether such an experience is possible. Note that there are two criteria that our account of her experience must satisfy: her experience must seem photograph-like during normal perception, but we must also account for how she is able to experience her flaws under non-normal conditions. One suggestion is that the functional flaws are “filled-in.”³² “Filling-in,” at least as it is understood with respect to this debate, refers to a theoretical process whereby the brain edits and adds information (hence “fills-in”) in the regions of the visual field that are missing visual information due to functional flaws—allegedly accounting for why we do not notice these gaps of missing information.³³ Suppose that we wish to entertain some variety of “filling-in.” This would allow for the Ordinary Perceiver to have the experience requisite for acquiring the belief. Would it also warrant us in attributing to her the Grand Illusion? It does not: her belief that her vision is everywhere detailed, distinct, continuous, in color, photograph-like, etc., would be *true*; her experience would not be illusory

experience failed to conform to her expectations as generated by an inference-based belief, it seems to be a misnomer to call her previous experience an illusion, in the same way that it seems incorrect to say that I suffered from an illusion in my “shirt collar” example.

³² I intend for this term to be understood, as used here, as referring to a theoretical process whereby the “gaps” that are a consequence of the flaws are touched up in such a way (the bit-map “filling in” of Dennett’s *Filling-in vs. Finding Out*) that the gaps are eliminated. We perceive the experience as “gap-free.”

³³ I contend that the argument that we ought to experience these “gaps of information” under normal circumstances (or, at least, that is what perception *really is like*) is invalid. This is because it does not follow from the fact that we can experience these flaws under certain circumstances that either (a) we ought to expect to experience them at all times, or (b) to not experience them under other circumstances is indicative of being subject to an illusion.

because the filling-in would make her vision seem that way because it *is* that way. In other words, according to filling-in, the perceptual product—the object of perception—is everywhere detailed, distinct, continuous, in color, etc., such that photograph-like perceptual experience is veridical. However, the challenge³⁴ for those who would argue for filling-in is to explain how it is possible for us to *ever* experience our flaws. Given that the nature of “filling-in” is such that that we do not experience the flaws because what we perceive is a filled-in perceptual product which is developed further down the causal chain (retinal image is touched up), it is not clear why filling in would fail us (since the image does not depend strictly on the incoming information—it is an edited version of the original information).

Perhaps an advocate of filling-in might object to my dismissal of her position as a viable explanation for the Ordinary Perceiver’s experience.³⁵ After all, doesn’t our failure to experience these flaws warrant *some* explanation—and what could be better than an explanation that argues that we would not expect to experience them because they are “filled-in”? We might inquire as to how the “filling-in” proceeds—that is, how it is that these flaws are edited. The advocate of a veridical variety of “filling-in” could answer, as so many have, that the brain “knows what goes there” and “fills it in” accordingly. However, this explanation has a bizarre, undesirable consequence: we would expect that, in unfamiliar environments, our functional flaws would initially be exposed until our brains figure out “what goes there.” Since we do *not* experience this occasional exposure of the flaws,³⁶ we can reject explanations that

³⁴ The need to meet such a challenge is removed once the Grand Illusion hypothesis is rejected in favor of my proposed reformulation of our expectations. See footnote 29, below.

³⁵ See my argument for how, once we have modified our expectations appropriately, the advocate of filling-in no longer owes any explanation regarding why filling-in “fails” when we perceive our flaws. Furthermore, all of the undesirable consequences (both here and elsewhere in the debate) are removed, such that filling-in can be seen as a viable theory of perception, rather than the oddity that it is taken to be at this point in the discussion.

³⁶ I can assert this with confidence, given that Ordinary Perceivers exhibit surprise upon experiencing their flaws, a reaction that we would not expect if they experienced them occasionally; it is the novelty

posit veridical filling-in³⁷

Suppose that one were to argue for a non-veridical variety of filling-in. The bizarre consequence of the veridical filling-in account would be avoided—but at what cost? One might argue that the content need not be veridical—any filling-in will do, just so long as the perceiver does not experience the gaps. However, this has a bizarre consequence of its own. Depending how this filling-in is selected, we might be *more* likely to notice our functional flaws: if the assignment of the filling-in design (what the filling-in looks like) is random, then we would expect our functional flaws to be obvious to us. The content of the filling-in could be different from the incoming visual information around it, such that we could more easily detect the flaws due to the anomalous information.³⁸ The advocate of non-veridical filling-in might counter by arguing that, since these flaws are not located in the foveal region, like the rest of the visual information that we acquire external to the fovea, we do not experience it in enough detail to notice the difference. However, this claim is not valid: while their assertion about not experiencing visual information acquired external to the foveal region as highly detailed is true, the same argument cannot be made for the contents of filling-in—unlike the visual information, which requires photoreceptors to be acquired, filling-in is argued to be generated by the brain, which is not subject to the same restrictions. One could, of course, argue that the brain represents that content as poorly as the visual information attained through that region external to the fovea (perhaps so as to not create disparities between the “filling-in” and the surrounding visual information), but this begins to look like an unnecessarily complex

of the experience which warrants the surprise (and which would be lacking if this explanation were true).

³⁷ At least insofar as the account of “filling-in” requires “knowing what goes there” (and, perhaps due to my unfortunate lack of imagination, I cannot fathom any other way of attaining veridical “filled-in” content).

³⁸ Rather than experience the flaws as poor or absent visual information, we would now experience them as containing anomalous information as compared to the rest of the visual field.

explanation.

Could our Ordinary Perceiver gain the experientially-based belief any other way? It seems that she could not. Her being able to acquire the belief through experiencing her vision as everywhere detailed, distinct, continuous, in color, photograph-like, etc., requires that her vision *is* this way (thus enabling her to experience it as such).³⁹ However, if her vision *is* this way, then she is not subject to any illusion—the content of her experience does not deviate from the predicted nature of experience⁴⁰ and her belief that her experience is “photograph-like” is true.⁴¹ Indeed, any experientially-based belief that is acquired through experiencing veridical, filled-in content will render (1) false; there is no discrepancy between the nature of our vision and how it seems to us. This rules out the experientially-based belief with the content ‘my vision is everywhere detailed, distinct, continuous, in color, etc.’ as a viable option.

2.4 OTHER CANDIDATES FOR THE BELIEF

In order to be justified in attributing the Grand Illusion to the Ordinary Perceiver, we must discover whether the belief that her vision is everywhere

³⁹ There is, perhaps, one other option, which I do not entertain here (because I do not see any reason to take it seriously): continuous hallucination. This would not require “filling-in” (this is different from the brain “touching-up” the contents, and could be entirely accidental) but would involve constant hallucination in the regions of those flaws.

⁴⁰ This is because the object of perception is, on such an account of filling-in, a touched-up image, which *is* everywhere detailed, distinct, continuous, in color, etc. Given that an account that uses filling-in as an explanatory resource must posit that perception is edited in this way, then our expectations would be altered from the original set in the debate—we would expect that our vision is everywhere detailed, distinct, continuous, in color, etc., because that which we perceive (the filled-in perceptual product) is that way.

⁴¹ Further, it is worth noting that it is unclear how we could *ever* experience the flaws (if they are “filled-in”). Perhaps the guise is only active when we are not attending to it or, like a poorly placed mask (think of enjoying a play from a seat where one’s angle to the stage is such that one only sees the sides of the masked actors—we would see what the person on stage *really* looked like (behind the mask), and not the mask—whereas the rest of the theatre enjoys the disguise), we are unable to see it under certain circumstances. Still, this begins to feel ad hoc—the more “just so” our account becomes, the more we may rightly long for a stronger explanation.

photograph-like, is the only belief that would cause her to exhibit such surprise and, if not, whether it is at least equally as strong of an explanation as the attribution other beliefs (regarding how well it fits the relevant information). As such, we must evaluate other candidates for the belief that would cause her to exhibit this surprise. If any of these candidates are viable, we must determine whether the belief_e that her vision seems photograph-like has at least comparable, if not greater, explanatory strength.

To illustrate that viable alternatives exist,⁴² we will address one such alternative raised by Jonathan Cohen.⁴³ He suggests that, rather than having theoretical commitments regarding the nature of her visual field (which he believes that we are not warranted in attributing to her), the Ordinary Perceiver simply believes that they notice objects and events that are located or occurring in front of them. Such a belief would also provoke surprise in light of experiencing the functional flaws—it would generate expectations for the Ordinary Perceiver that (it seems to her as though) she never fails to notice objects and events that are directly in front of her. The Ordinary Perceiver would be surprised upon discovering her blind-spot (where part of the object(s) are occluded), experience change-blindness or inattentional-blindness (failing to notice events happening before them due to lack of attention) because these experiences fail to conform to the expectations that would be caused by such a belief.

As such, the requisite belief is not the only one which would elicit surprise in

⁴² Due to the length of this paper, I omit my discussion of other potential explanations of why the Ordinary Perceiver is surprised by experiencing her functional flaws, such as other candidates for the belief. The importance of illustrating that there is at least one other candidate for the belief serves as evidence that we must continue our inquiry; that the Ordinary Perceiver is surprised—and that the surprise *can* be explained by the requisite belief—is insufficient grounds for attributing to her the Grand Illusion. However, see also my later discussion of beliefs about the nature of occlusion as another potential explanation.

⁴³ Cohen, *The Grand Grand Illusion Illusion*

the Ordinary Perceiver. Since this is so, we must further examine the beliefs that we could attribute to the Ordinary Perceiver to determine which is the best explanation of her surprise (and whether the surprise is indicative of the Grand Illusion).

2.5 FURTHER CONSIDERATION OF INFERENCE-BASED BELIEFS

If I am correct in concluding thus,⁴⁴ then the only candidate (with the content ‘my vision is everywhere detailed, distinct, continuous, in color, etc.’) for the cause of the Ordinary Perceiver’s surprise is a belief_i. This belief_i would cause her to exhibit surprise upon discovering the functional flaws of her visual system. The Ordinary Perceiver would have the tacit belief⁴⁵ that it is not the case that her vision seems to her as though it is not everywhere detailed, distinct, continuous, in color, photograph-like, etc. One might wonder why I assert that beliefs_i are held *tacitly*. Arguably, not all our beliefs are held overtly; some of our beliefs, though we may not actively entertain them or have the ability to express them linguistically (or even know that we have these beliefs), nonetheless impact our behavior. For example, though an Ordinary Speaker of the English language may not have explicit access to her beliefs about the conjugations of words in the English language (and, given what I know about the teaching of the English language in most schools, it is safe to assume that most Ordinary Speakers of English do not), she can nonetheless detect the occurrence of a conjugation error. This is similar to what Pylyshyn refers to as ‘tacit knowledge’: though the individual might not be able to make use of such information for the purposes of answering questions about the matter, she can reveal such knowledge by

⁴⁴ That there is no viable account that would explain how she could experientially-acquire the requisite belief while nonetheless still be able to experience her flaws under non-normal perceptual circumstances.

⁴⁵ By ‘tacit belief,’ I mean roughly what Pylyshyn describes as ‘tacit knowledge,’ though I do not intend to imply any further similarities between the ‘mental imagery debate’ and the Grand Illusion at this point.

exhibiting telling behavior when occurrences are contrary to her knowledge.⁴⁶ Likewise, this would account for both the surprise exhibited by the Ordinary Perceiver upon experiencing the flaws *and* her strange looks and responses when queried about whether she believes that their vision is everywhere detailed, distinct, continuous, in color, photograph-like, etc.⁴⁷

Perhaps one might object that my explanation is incomplete yet; I have not yet shown how inference-based beliefs do not warrant our attribution of the Grand Illusion to the Ordinary Perceiver. After all, it appears that the content of the inference-based belief is at odds with the “real” nature of perception.

To explain how inference-based beliefs do not warrant the attribution of the Grand Illusion to the Ordinary Perceiver, it is useful explicating two seemingly disparate accounts by Noë and Dennett, illustrating how these accounts can be interpreted as sharing much common ground. This common ground illustrates that the surprise, if generated by inference-based beliefs, is insufficient to warrant attribution of the Grand Illusion to the Ordinary Perceiver.

In an attempt to explain why the Ordinary Perceiver is surprised by her functional flaws, Noë introduces the problem of *perceptual presence*.⁴⁸ This problem is a strengthened version of the Grand Illusion hypothesis. The worry is that we seem to “experience” a whole object when we are, in fact, only perceiving part of it. For example, when we close our eyes and hold a bottle in our hand, it seems to us as though we are experiencing the whole bottle when, in fact, our hand is only making contact with a portion of it. Likewise, when we see a cat on the other side of a fence, though it seems to us that we experience the whole cat, we only perceive part of it

⁴⁶ Pylyshyn, *Mental Imagery: In Search of a Theory* (particularly section 3.1)

⁴⁷ Ordinary Perceivers generally reply to inquiries about their visual experience with either puzzled looks or thoughtful replies stating that either they do not believe their experience to be that way or that they had never thought about it.

⁴⁸ Noë, *Is the World a Grand Illusion?*

(that which is visible between the boards of the fence).

Noë argues that the key to resolving the problem of perceptual presence is through analysis of amodal perception. Amodal perception is, “perceiving what is, strictly speaking, out of view.”⁴⁹ Just as we amodally perceive the hidden parts of the cat, Noë suggests that we amodally perceive objects in our environment which we are not currently *visually* experiencing. He states that this proposed alteration to the problem changes the question from one of whether the Ordinary Perceiver has the (illusory) experience as though she perceives everything in the environment at once to one of whether amodal experience is illusory. Noë argues that there is no need for the brain to create a representation of, say, the whole cat (which is partially occluded by the fence) since all we need to do to get more information about the cat is to move around so that different parts of the cat become visible. We are skilled perceivers; we know how to manipulate objects and adjust ourselves such that we are able to perceive the information that we seek. When we squint or peer, as Noë puts it, we do so because that which we wish to see is yet too far or too small from our current visual perspective. We know what our range of visual clarity is and we adjust accordingly. Too big? Back up. Too small? Lean in.

At the end of the day, though Noë and Dennett have argued against one another’s positions regarding whether the Ordinary Perceiver is subject to the Grand Illusion, it appears that they have much in common. Both argue that the Ordinary Perceiver does not *experience*⁵⁰ vision that seems everywhere detailed, distinct, continuous, in color, photograph-like, etc.; rather, their experience is a matter of

⁴⁹ Noë, *Is the World a Grand Illusion?*

⁵⁰ Though Dennett claims that it does *seem* that way to the Ordinary Perceiver—contrast with Noë, who argues that it does not seem that way to her (the phenomenology involved in the Grand Illusion hypothesis is in error with respect to how it seems to her). Noë argues that, rather than perceiving a line which extends through the region of your blind spot (while you are attending to that region) as being unbroken, it seems to her as through her experience of the line is that it does not seem to be broken

representation, interpretation, or skilled (amodal) perception of those regions that fall outside of the region to which the perceiver is currently visually attending. Both accounts rely on a judgment sense of “seeming” (amodal perception, “finding out”), not on the experiential sense of the term. Understood in this way, both of their positions are consistent with arguing that the relevant belief held by the Ordinary Perceiver is an inferentially-based⁵¹ belief. The ‘amodal perception explanation’ offered by Noë appears to be very similar to the account that Dennett gives as an explanation of our experience of Andy Warhol’s Marilyn Monroe wallpaper—both argue that we do not *experientially* represent the poor or absent information as being present. Dennett writes:

“Suppose you walk into a room and notice that the wallpaper is a regular array of hundreds of identical sailboats, or—let’s pay homage to Andy—portraits of Marilyn Monroe. We know that you don’t foveate, and don’t have to foveate, each of the identical images in order to see the wallpaper as hundreds of identical images of Marilyn Monroe. Your foveal vision identifies one or a few of these and somehow your visual system just generalizes—arrives at the conclusion that the rest is “more of the same.” We know that the images of Marilyn that never get examined by foveal vision cannot be identified by parafoveal vision—it simply lacks the resolution to distinguish Marilyn from various Marilyn-shaped blobs. Nevertheless, what you see is not wallpaper of Marilyn in the middle surrounded by various indistinct Marilyn-shaped blobs; what you see is wallpaper composed of hundreds of identical Marilyn. Now it is

⁵¹ Presumably, Noë would not use the term “inference.” However, I only intend to imply that both his account and Dennett’s “finding out” involve something like a judgment. For example, on both accounts the belief that vision seems “photograph-like” is a judgment about the nature of vision based on the parts of the visual field that the perceiver can experience (and, from that, she judges that the flaws contain similar visual information), and not because she *experiences* the flaws as being “filled-in”, in the sense that there is some perceptual product which she *experiences* as “photograph-like”.

a virtual certainty that nowhere in the brain is there a representation of the wall that has high-resolution bit-maps that reproduce, xerox-wise, the high-resolution image of Marilyn that you have foveated. The brain certainly would not go to the trouble of doing that filling in! Having identified a single Marilyn, and having received no information to the effect that the other blobs are not Marylins, it jumps to the conclusion that the rest are Marylins, and labels the whole region "more Marylins" without any further rendering of Marilyn at all. Of course it does not seem that way to you. It seems to you as if you are actually seeing hundreds of identical Marylins. And in one sense you are: there are, as I said, hundreds of identical Marylins out there on the wall, and you're seeing them. What is not the case, however, is that there are hundreds of identical Marylins represented in your brain. Your brain just represents that there are hundreds of identical Marylins."⁵²

Notice another similarity between their accounts: they both offer interpretations that fit the behavior of the Ordinary Perceiver. Both accounts are consistent with the surprise (due to experiences that are contrary to her expectations) and lack thereof (due to experiences that conform to her expectations) that is exhibited by the Ordinary Perceiver. Neither account attributes to her expectations (or beliefs) that would lead her to fail to be surprised when experiencing the flaws—her surprise is consistent with both accounts, and the same is true for her lack of surprise. Thus, the assertions of each regarding what does and does not surprise the Ordinary Perceiver fit equally as well with the other's account as their own.⁵³

I will now consider whether the content of the inference-based belief is such

⁵² Dennett, *Filling In versus Finding Out*

⁵³ This is particularly interesting when one considers that the role that Dennett's assertions regarding what *does* surprise the Ordinary Perceiver and Noë's assertions about what does *not* surprise the Ordinary Perceiver each appear to be intended as evidence against the other's account. However, in this light, it appears that both sets of assertions fit *each* account equally well.

that we are warranted in attributing the Grand Illusion to the Ordinary Perceiver. Recall that the inference-based belief in question is held tacitly. If the Ordinary Perceiver tacitly believes that it is not the case that her vision seems to her as though it is not everywhere photograph-like then her belief does not entail that she also believes that her vision *is* everywhere detailed, distinct, continuous, in color, photograph-like, etc. No beliefs that assert positive claims are made about how her vision seems to her—only a negative one: it isn't the case that her vision is not photograph-like. She could acquire this negative belief because she has no reason to accept its negation: that her vision is not everywhere photograph-like. In other words, in order for her to acquire the positive belief that her vision is not everywhere photograph-like, she must have an experience that causes her to have that belief (such as experiencing the flaws). In the absence of such experiences (and experiences that would cause the belief that her vision is everywhere photograph-like—such as those discussed above), she is committed to neither belief and instead acquires the tacit belief that it is not the case that her vision seems to her as though it is not everywhere photograph-like. This tacit belief roughly amounts to the expectation discussed by Cohen: the Ordinary Perceiver believes that she perceives the events and objects that are directly before her.⁵⁴ In other words, her expectation is governed by her tacitly held belief that it is not the case that her experience is not everywhere photograph-like (she does not expect that she will fail to experience some of the objects and events immediately before her).

2.6 WHY DENNETT CANNOT ARGUE THAT THE ORDINARY PERCEIVER IS SUBJECT TO THE GRAND ILLUSION

Recall my earlier argument that the belief which is allegedly overthrown when

⁵⁴ Cohen, *The Grand Grand Illusion Illusion*

the Ordinary Perceiver is surprised at the experience of her flaws must satisfy two criteria: (a) it must have the content that vision seems to be everywhere “photograph-like,” and (b) this belief must be acquired through experience (her vision must seem “photograph-like,” in the experiential sense of seeming). We have seen that one cannot argue for the Grand Illusion by using the surprise as evidence unless the belief in question was acquired through experience. Now we will look at why one cannot make the case for the Grand Illusion unless one argues that beliefs have determinate content.

As noted earlier, Noë argues that Dennett is the main proponent of the Grand Illusion Hypothesis. However, I contend that he cannot argue for the Grand Illusion in the way that he attempts to do so. Since he is taken to be a central figure in the debate, it is worth reviewing the problem with his account which prevent him from being able to argue for the Grand Illusion, as well as one regarding his explanation for why the Ordinary Perceiver fails to experience her flaws.

There are two problems with Dennett’s account of the Grand Illusion hypothesis: the first is regarding how he attempts to establish his case for the Grand Illusion, and the second is regarding whether his “labeling” violates his Thrifty Producer Principle. Since Dennett’s account of how he attempt to establish his case for the Grand Illusion hypothesis is similar to other accounts that we have discussed which support their claim by noting the surprise exhibited by the Ordinary Perceiver, we are already familiar with that strategy. However, Dennett’s case is special, in that his holism appears to prevent him from being able to establish the Grand Illusion hypothesis. Once we have seen why this is so, we may then inquire as to whether one aspect of his account regarding why we do not experience our flaws violates a central principle in his general position.

Dennett argues that the surprise exhibited by the Ordinary Perceiver when she first experiences her flaws is indicative of a belief being overthrown—a belief that her vision is everywhere detailed, distinct, continuous, in color, etc. However, Dennett, unlike other advocates of the Grand Illusion hypothesis, does not mean that she has a belief with such determinate content; as a holist, he argues that beliefs have indeterminate content. It seems consistent with his overall position to interpret him as attributing this belief to her as part of an intentional stance that he adopts towards the Ordinary Perceiver. As such, the belief is attributed as a means of explaining her behavior.

This seems right so far as it goes—if she had such a belief, *then* we would expect her to exhibit surprise when she first experiences her flaws. The interesting question is whether there is reason to think this is the *best* stance to adopt towards her behavior. Important considerations regarding whether this is the best explanation⁵⁵ include how well it integrates with the rest of the beliefs that we attribute to her when we have adopted stance with the strongest explanatory power. In other words, are there beliefs that we attribute to her which would have stronger explanatory power regarding her behavior, as well as fit better with the rest of the best stance that we adopt towards her?

Compare this belief to one regarding the nature of occlusion. For example, it seems that we could explain her surprise does not indicate when she discovers her blind spot by attributing to her a belief about the nature of occlusion (as supported by her experiences thus far)—namely, the belief that an object (or some part thereof) is occluded from view if and only if it is blocked by another object. Surely, to one who

⁵⁵ Since his Intentional Stance is a predictive and explanatory strategy, meaning that we adopt a particular stance both to make sense of previous behavior and as a means of predicting future behavior (by hypothesizing how an ideally rational system might behave).

holds such a belief, it would seem surprising for an object to be occluded from view without being hidden behind another object. However, such surprise does not indicate that she is subject to the Grand Illusion; though her experience in this case fails to conform to her expectations, it is this, and not an illusory experience, that generates the surprise. As such, it appears that there are other beliefs that we can attribute to her which explain the surprise but which fail to indicate that she is subject to illusion.

Regarding which belief we should attribute to her, we need to consider which best fits with the rest of the intentional stance that we adopt towards her. While the former belief fails to connect up with any other beliefs we can attribute to her (perhaps this is why we are so resistant to saying that the Ordinary Perceiver really does have such a belief—it is so unlike the other beliefs that we could attribute to her that we have a hard time attributing such a belief), it appears that the latter could easily fit within the framework of cognitive development. For instance, if we assume that she acquired her belief about the nature of occlusion through previous experience, namely by having noticed that the only time that objects disappear is when they are hidden behind other objects, then she would not expect for an object to be occluded in that region of visual field (given that there is no object there which would occlude it).

Regardless of whether we can determine which belief best fits with intentional stance that has the greatest explanatory and predictive power, if attributing a belief with this content to the Ordinary Perceiver is necessary for making the case that she is subject to the Grand Illusion, then we ultimately end up with a skeptical problem. To understand how this problem arises we must first get clearer on how, exactly, Dennett attempts to make his case. Once we see this, it will become apparent that, in light of his holism, Dennett's position entails that there is no fact of the matter whether the

Ordinary Perceiver is subject to the Grand Illusion.

Recall that the importance of belief is merely to reveal the content of the Ordinary Perceiver's visual experience. When she exhibits surprise, we can infer that her experience failed to conform to the expectations generated by her beliefs. Thus, the surprise she exhibits upon first discovering her flaws is taken to indicate that a belief with the contents which would generate expectations which are contrary to her experience of flaws, which was generated by past experience, is overthrown.

Holism, however, entails that no beliefs have determinate content; the content of beliefs are indeterminate, meaning that there is no fact of the matter what is the content of belief. This is problematic because, as we saw earlier, the belief with requisite contents is necessary for establishing the Grand Illusion hypothesis in this way. Regardless of which belief we attribute to her, Dennett must⁵⁶ conclude that there is no determinate content to her belief, such that it is indeterminate whether she is subject to such illusion.⁵⁷

Regarding the second problem with his account, Dennett's explanation for why we fail to experience our functional flaws has two parts: (1) the brain ignores the missing data and (2) the brain "finds out" what information falls within the areas of missing data, and labels that region accordingly. Central to his argument that we ought to reject filling-in in favor of his account is his Thrifty Producer Principle—"if there's no one to look at it, don't bother making it". In other words, if there is no part of the brain designated to processing such information, it is superfluous for such information to be generated (presumably by whatever mechanism would be

⁵⁶ In light of his holism.

⁵⁷ Note that I am not arguing about whether beliefs have determinate content. I am only arguing that, given the use of the surprise as the means of determining whether the Ordinary Perceiver is subject to the Grand Illusion, the belief in question must have determinate content to make the case in this way. But this, of course, is independent of my own position regarding whether beliefs are determinate, which is irrelevant here.

responsible for the filled-in information). However, this principle in conjunction with the two parts of his explanation for why we fail to notice her flaws makes finding out superfluous as well as a violation of principle.

His account begins with arguing the brain simply ignores the missing data. He maintains that there is no part of the brain dedicated to processing visual information from the regions the visual field which, due to the perceptual flaws, lack such information to some extent. Since there is no part of the brain assigned to the task, the missing information goes unnoticed.

It is this claim that there is no part of brain assigned to the task of processing what would be the visual information from the regions which contain the flaws which motivates his Thrifty Producer Principle; that there is no part of the brain assigned to task seems (to Dennett, anyway) to entail that there is no part of the brain which would perceive the touched-up perceptual product, making any such product superfluous. Therefore, he concludes that there must not be any filling-in.

He then goes on to argue that the brain “find out” what falls under those regions of the visual field which are missing information due to the flaws and labels that region accordingly. However, this part of the account is problematic for two reasons. The first is that it is superfluous, given that he has already argued that the brain ignores the missing data. Secondly, it violates the thrifty producer principle.

An example of how “finding out” works will be useful here. Recall the earlier quoted passage in which he illustrates this point through asking the reader to imagine that she is standing before a wall full of Marilyn Monroe image wall paper by Andy Warhol. He argues that, though it seems to the Ordinary Perceiver as though she is seeing hundreds of Marilyn images, is not the case that, somewhere in her brain, there is a representation of hundreds of Marilyn images; rather, from the images that she

does see before her, and no disconfirming information to the contrary from the periphery, the brain just “jumps to the conclusion” that it is more of the same, and labels that region “more Marilyn images”.

Finding out is unnecessary given that he has already argued that the brain ignores the missing data. If the missing information is ignored, then it is not clear why we need to posit the labeling process. More specifically, it is not clear what work the labeling is doing. The labeling appears to be intended to play an explanatory role in his account regarding why we do not notice the flaws, but the ignoring feature of his account appears to fully explain this independently of the labeling. Furthermore, it is not clear how we could ever experience the labeling, since this seems like a perceptual product, and Dennett has argued that there is no part of the brain dedicated to perceiving such products.

Also problematic is that labeling appears to be at odds with the principle. As we have just seen, he argues that there is no perceptual product. However, labeling does appear to be just such a perceptual product. It is unclear what work labeling could do if we fail to have the mechanisms necessary to experience it.

Give these considerations, it seems that the best way to preserve the most of his account is to simply reject the labeling feature. When we exclude the labeling feature, the account is no longer internally incoherent. The ignoring feature and the principle are consistent with one another. Furthermore, this allows him to maintain his Thrifty Producer Principle, which plays an important role in his attack on the Cartesian Theatre conception of mind.⁵⁸

⁵⁸ This is one of the central moves of his position in general—he has been a long-standing opponent of such a theory of mind. (See ‘Consciousness Explained’, for example)

3. REEVALUATING THE ASSUMPTIONS THAT MOTIVATE THE GRAND ILLUSION HYPOTHESIS

3.1 WHAT ASSUMPTIONS MOTIVATE THE HYPOTHESIS?

Since our attempt to give an account regarding how it is that the Ordinary Perceiver's vision seems photograph-like under normal circumstances, but allows her to experience her flaws under special circumstances, seems fruitless, it will be useful to critically review the steps that lead us to this point. It seems that if we could somehow show that our conception of predicted (non-illusory) nature of perception was such that, if our expectations were different (that is, such that we wouldn't expect to experience the flaws under normal conditions), then we would no longer have the burden of explaining why we don't experience them normally but do experience them under certain circumstances. In other words, since the Grand Illusion hypothesis is motivated by the expectations generated by our assumptions (regarding the affect that the flaws would have on our perceptual experience), by reevaluating these assumptions we may find a way of altering our expectations, such that we no longer need to offer an account as to how this illusory experience of failing to experience the flaws is possible.

Recall that Dennett assumes that the surprise exhibited by the Ordinary Perceiver when she discovers her flaws indicates that she was expecting that her experience would be different from that what she actually experienced in this case. From this, he appears to assume that, since she was surprised by her experience of the flaws (the experience of which fails to be photograph-like), then she was expecting her experience to seem photograph-like, and she had such expectations because her previous experience had seemed photograph-like to her. However, this conclusion does not follow. It seems that it is reasonable to infer that her surprise only indicates

that she was not expecting to experience what she had when she discovered a flaw. In other words, since she was surprised to experience her vision as failing to seem photograph-like, all we can infer is that she was not expecting for her experience to not seem photograph-like, but from that it does not follow that she did expect to experience to seem photograph-like.

One might wonder why I argue that Dennett implicitly assumes (or, at the very least, *inadvertently* assumes, such that it is entailed by his position) that she is expecting that her experience will be photograph-like. The following seems to be a reconstruction of how he determines her expectation. Consider the role that the surprise performs in determining whether the Ordinary Perceiver is subject to the Grand Illusion. Since she is surprised to discover her flaws, Dennett infers that she was expecting something else, namely that her experience would be photograph-like. This expectation must have been caused by a belief that her experience would be so—and this belief must have been generated by previous experience(s) which did seem photograph-like to her (at the time that the belief was formed and throughout the duration which she held the belief). He happily notes that her experience did seem photograph-like to her, and he concludes this by taking as evidence the surprise that she exhibits when she experiences her flaws for the first time. The only way that his conclusion (that her previous experience seemed photograph-like to her) follows from the surprise is with the inclusion of the missing premise that the overthrown belief is that her vision seems photograph-like, such that when she experiences her vision as not-photograph-like (during the discovery of her flaws), it is *that* belief that is overthrown. Otherwise, all he would be warranted in concluding from the surprise is that she was not expecting the not-photograph-like experience, but this belief could be generated by more than just the photograph-like experience. Whereas, the only belief

(overthrown upon discovery of the flaws) that would support his claim that her experience seemed photograph-like would be the experientially-acquired belief that her vision seems photograph-like, such that the “something else” that she was expecting was that her experience of the region of, say, her blind spot, would not be not-photograph-like (since she experiences that area during normal perception as photograph-like).

The surprise exhibited by the Ordinary Perceiver could be interpreted as indicative of both (by which I mean her expectation that her experience would be photograph-like or that it was not the case that she expected her experience to not be photograph-like would both cause surprise). However, in the case that she was not expecting that her experience would not be photograph-like, the Ordinary Perceiver cannot be said to suffer from the Grand Illusion, because we can only establish that it is not the case that her vision seems to her as though it is not everywhere photograph-like. But this is different from seeming to the Ordinary Perceiver that her vision *is* everywhere photograph-like. These beliefs are functionally equivalent (they cause the Ordinary Perceiver to behave in the same way—at least regarding the display of surprise), but their content is distinct. However, from the fact that it does not seem to the Ordinary Perceiver that her vision is not everywhere detailed, distinct, continuous, in color, etc., it does not follow that her vision does seem to her to be everywhere detailed, distinct, continuous, in color, etc. Her vision could appear many different ways and still satisfy the criterion ‘does not seem to her as though her vision is not everywhere detailed, distinct, continuous, in color, etc.’ Only if we can make the case that she has the belief_c that her vision seems everywhere detailed, distinct, continuous, in color, etc., are we warranted in inferring that it does seem to the perceiver as though her vision is everywhere photograph-like.

Consider this analogous case: our perceiver is walking through her local campus, when she gets hit with a water balloon. She is surprised by this, but why? Is it because she is expecting to not get hit with a water balloon? Presumably not--she could be surprised by this event even if she has no previous experience with water balloons (such that she is not aware that such things even exist). If surprise was only possible in cases where one was expecting the exact opposite event from the one experienced (in this case, if the perceiver could only be surprised if she was expecting to not get hit with a water balloon), then the only circumstances under which one could be surprised would be those which one is both aware of the source of the surprise (minimally, that they know such a thing exists) and expecting that their experience will be such that the source of the surprise will be absent. However, this entails that, by definition, the Ordinary Perceiver is not surprised by her experience of the flaws. This is because, due to the criterion of being unacquainted with the flaws as a prerequisite for being Ordinary Perceiver, she does not know that the source of her surprise exists, such that she could not possibly expect not to experience it.

One might protest that the advocates of the Grand Illusion hypothesis never argue that she was expecting the exact opposite of her experience and, as such, my claim that it is entailed by their position is unjustified. While it is true that no party explicitly argues as such, my conclusion is still entailed by their position. One can see this by looking at how they interpret the surprise. From the Ordinary Perceiver's surprise, they argue that she was expecting something else—which ultimately turns out to be a continuation of her previous, "photograph-like" experience. What is of interest here is how they get from her surprise to the conclusion about the "photograph-like" quality of her experience. Since she was surprised because her experience did not conform to her expectations (which are generated by a belief that is

overthrown when she experiences her flaws for the first time), it appears to be implicitly assumed that her belief was caused by her experience being "photograph-like" due to the widespread discussion regarding why it is that she fails to experience her flaws. To put the point another way, it is unclear why there would be so much debate over how we should account for her failure to experience the flaws if it is not assumed that, under normal conditions, she perceives the area which contains the flaws (but nonetheless fails to experience them). And, of course, I am not arguing that the Ordinary Perceiver analyzes her experience in this way. But this is in contrast to the belief that is requisite for warranting our attribution of the Grand Illusion to the Ordinary Perceiver. If it is not the case that her vision seems to her as though it is not everywhere detailed, distinct, continuous, in color, photograph-like, etc., this is different from seeming to her that her vision *is* everywhere detailed, distinct, continuous, in color, photograph-like, etc. The inference-based belief does not include any positive assertion about what her experience *is* like (only what it is *not* like); numerous different positive assertions about the nature of her experience are consistent with this tacit, inference-based belief. Since this belief does not have the requisite content, we cannot attribute the Grand Illusion to her. The onus is on the advocate of the Grand Illusion to give an account of how the Ordinary Perceiver can experientially acquire the belief that her vision is photograph-like that does not run into the problems already raised.

Given that these expectations are what motivate the claim that the Ordinary Perceiver is subject to an illusion, it seems plausible that by changing our assumptions, we may be able to eradicate the problem. To figure out how to change our expectations, we must carefully determine what they are, and see what assumptions led to these expectations. Once we have done this, we can challenge

those assumptions to determine whether we ought to adopt new assumptions which would alter our expectations in the desired way.

3.2 A DISTINCTION BETWEEN KINDS OF FUNCTIONAL FLAWS

It seems that the best place to start our evaluation of the assumptions that motivate the Grand Illusion hypothesis is right back at the beginning, with the functional flaws. As noted earlier, there are many different flaws, the failure to experience all of which is taken to be indicative of the Grand Illusion. However, we can make a useful distinction between two kinds of flaws based on the features of the visual system which generate them.

As noted above, part of our explanation must include attention to the predicted nature of perception. I believe that the predicted nature of perception, as understood in the Grand Illusion, rests on an invalid argument. Recall that, by the predicted nature of experience (as formulated within the Grand Illusion hypothesis), I mean the nature of perception as including the experience of the functional flaws. Given that it is assumed that perception that fails to include experience of the flaws is illusory, then presumably, if our experience included the flaws, we would not be subject to an illusion. This is derived from an argument like the following:

4. We have functional flaws, which we can experience under certain conditions.
5. We do not experience these flaws under normal conditions.
6. Since the Ordinary Perceiver has not yet experienced these flaws (and, upon doing so for the first time, exhibits surprise), we may infer that she was expecting something else to happen—thus overthrowing her belief that caused her expectation to the contrary (that her vision is detailed, distinct, continuous, in color, photograph-like, etc.). She is subject to an illusion—she does not experience her

vision as it *really is*.

However, it does not follow from the fact that we have these functional flaws which we are able to experience under certain conditions that we must notice them under other or all conditions. Furthermore, it does not follow that the predicted nature of experience which motivates our expectation that we would experience our flaws under normal conditions is based on the best possible set of assumptions that could be formulated based on the data from the relevant perceptual research.

This appears to be a problem with our expectations: because we *know* that the flaws exist, we *expect* that we should experience them. However, it is not clear that we are warranted in having such expectations. As noted above, from the fact that we have these flaws which we can experience under certain conditions, it does not follow that we ought to experience them in any other or all circumstances.

As we saw, the Grand Illusion hypothesis entails an incoherent formulation of the predicted nature of perception. Though never explicitly stated, this formulation can be directly derived from that which is taken to be illusory about the Ordinary Perceiver's visual experience. The error ultimately stems from considering the data from the two different kinds of flaws independently, then formulating expectations for each kind without considering the data of the other. When we simultaneously evaluate the two contrary accounts of predicted nature of perception, it is apparent that the formulation is incoherent.

A distinction can be made between two kinds of functional flaws. The differentiating feature between the two kinds is the aspect of the perceptual system that causes them. One kind is generated by the functional structure of the visual system. Let us call these the 'structural flaws'. The other kind is generated by the attention-dependent nature of perception. We will call these 'scope flaws', as these

flaws restrict the scope of that which we notice about the changing world before our eyes.

The structural flaws are caused by the functional structure of our visual system. By functional structure, I mean how the structure of our visual system impacts the quality of our visual information. Perhaps the best example of this kind of flaw is blind spot. The blind spot is created by the lack of photoreceptors in front of where the optic nerve connects to the back of the eye. Since there are no photoreceptors, we fail to receive visual information from the region.

Other examples of the structural flaws include the saccades, uneven distribution of rod and cones (rods primarily external to foveal region, cones restricted to foveal region) and a lessening in general of photoreceptors from the center of the visual field (where there is a high density of cones) towards the periphery (where there are no cones and very few rods).

It is important to understand the impact of the structural flaws both on her visual experience under normal conditions and also under contrived conditions which do not conform with our normal perception. The discrepancy between how the flaws impact our vision under normal circumstances and under special conditions leads the advocates of the Grand Illusion hypothesis to argue that our experience under normal circumstances is illusory. The assumption appears to be that, since we know that we can experience these flaws, we must be subject to illusion when we fail to experience them. In other words, since the functional structure of our visual system is responsible for her visual experience, then one must be subject to illusion if one's visual experience fails to conform to that which the functional structure would support. Thus, it is useful to discuss the effect each flaw has on our visual experience under both kinds of circumstances, according to the Grand Illusion hypothesis.

As noted earlier, the blind spot results from a lack of photoreceptors. Under normal conditions, we fail to notice our blind spot. However, we can easily experience our blind spot under contrived conditions, such as shutting one eye while fixing the gaze of the other on some fixed point, then attending to that region of the visual field (this usually requires some sort of image or symbol which we can observe to be altered by the blind spot—disappearing behind blind spot, etc.).

The uneven distribution of rods and cones results in differing quality of visual information from different areas of the retina. Cones are the photoreceptors that allow for color vision. They only work well in well-lit environments; they are functionally impaired in poor lighting. The high density of cones in the foveal region creates a highly detailed, richly colored visual experience of the center of our visual field. However, we have few rods in the area, making it very difficult to acquire visual information from that region in low lighting conditions.⁵⁹ Outside of the foveal region, we fail to have color vision. This is because rods, which are responsible for our ability to see under poor lighting conditions, do not support color vision.

Given the uneven distribution of, and differences in kind of visual information processed between, the two kinds of photoreceptors, the advocates of the Grand Illusion hypothesis conclude that our visual experience under normal conditions, which we do not experience as lacking color and detail in areas, must be illusory. Note that we are able to experience this kind of structural flaw under contrived circumstances. For example, we cannot detect the color, and sometimes even the shape, of an object that is held at our periphery. At best, we can detect it if it is moving—otherwise, we seem to be unable to discriminate anything about it.

Unlike the functional flaws, the scope flaws are not generated by the physical

⁵⁹ This is why, for instance, you cannot see most stars by looking directly at them. Instead, you must look slightly to the side of them.

structure of the visual system; they result from the attention-dependent nature of perception. Recent research has indicated that our visual experience is restricted to that to which we are currently attending. Two examples of scope flaws are change blindness and inattentional blindness. Both of these flaws are regarding our failure to notice large-scale changes and events that are occurring immediately before us. In both kinds of cases, we are unaware that we are failing to perceive any such changes and events.

3.3 FORMULATING THE PREDICTED NATURE OF EXPERIENCE

Independently of whether the advocates of the Grand Illusion hypothesis are warranted in their assumptions regarding what is indicated by the Ordinary Perceiver's surprise, let us grant them these assumptions for the sake of argument, so that we may see what these assumptions entail regarding the predicted nature of perception. To determine what is entailed, we must first formulate the predicted nature of perception with respect to each kind of flaw. We must then compare the two formulations, such that we can then see the entire account of predicted nature of experience as entailed by the assumptions of Grand Illusion hypothesis.

The Grand Illusion hypothesis appears to assume that, given what we know of the structural flaws, any visual experience which seems everywhere detailed, distinct, continuous, in color, etc., must be illusory. Suppose this were true—what does this entail about the predicted nature of perception? Since any visual experience that is “photograph-like” is illusory, and the claim that such experience is illusory is based on our expectations of how these flaws would impact our visual experience, then it follows that the predicted nature of perception would include experience of the structural flaws. For example, since our experience is said to be illusory based on the fact that we do not generally experience a lack of visual information resulting from

the lack of photoreceptors which is responsible for our blind spot, it follows that the predicted nature of experience would include such a lack of visual information.

As such, the predicted nature of perception would be highly detailed and richly colored only in a relatively small portion of the center of our visual field, with detail greatly decreasing towards the periphery. We would experience a lack of visual information in the region of the blind spot. Furthermore, our visual experience with fail to be continuous; due to the saccades, we would experience a rapid sequence of alternating black-and-white and color “snapshots”.

Regarding the scope flaws, the illusory experience is that it seems to us as though we perceive all of the changes and events that happen before us. However, as both kinds of scope flaws demonstrate, we do not notice all of them; what we perceive is constrained by that to which we are attending. Since the guiding assumption here is that it seems to us that we notice so much more of what's going on around us than we actually do, then I suppose that the predicted nature of experience with the such that it would be apparent that we only perceive that to which we are directly attending. Perhaps predicted nature of experience might be contained to some portion of the foveal region, such that the rest would be absent.

It is worth noting that here, as elsewhere, there is an equivocation on “seems”. One could mean that our visual experience is such that it seems to us as though we are noticing much more than we in fact are. This formulation seems strange to me—I am not sure what it would be like for my visual experience to seem to me as though I am experiencing much more of what is going on before me than I actually am. Maybe what is meant here is that the Ordinary Perceiver, who, for obvious reasons, does not notice that they fail to notice some of the events happen before them, forms the

implicit belief that they notice (all of the) changes going on in front of them.⁶⁰ We can then explain the surprise by appealing to that belief. However, this won't do either—the Grand Illusion is a *perceptual* illusion, not a false belief. The false belief, if not generated through illusory experience, is not indicative of a perceptual illusion.⁶¹

Note that all of these assumptions about the predicted nature of perception are derived directly from expectations generated by the research into these functional flaws. What is interesting is that the two kinds of flaws are never discussed together. They are either discussed in different sections of the same paper, or discussed in different papers altogether. As such, we can assume that the assumptions for each kind of flaw, regarding why failure to experience them is illusory, are independent. This is so because the case for each kind of flaw counting as evidence for the Grand Illusion hypothesis is made independently of any similar arguments for the other kind of flaw. Since the case for each kind of flaw is made independently, this entails that each kind of flaw has a distinct formulation of the predicted nature of perception based on the expectations generated by the assumptions regarding how that kind of flaw would affect our visual experience.

When we consider these assumptions together, our account of the predicted nature of perception becomes incoherent. The structural flaw formulation entails that we experience a lessening of color and detail towards the periphery, alternative sequences of black-and-white and color “snapshots”, no color outside the foveal region, etc., which requires the ability to perceive all of these areas of our visual field at once—which is exactly what the scope flaw formulation prevents by restricting our

⁶⁰ This is the kind of banal surprise that Jonathan Cohen attributes to the ordinary perceiver at the end of *The Grand Grand Illusion Illusion*.

⁶¹ See my earlier discussion of why this is so. Also, while it is interesting question regarding how it is that the Ordinary Perceiver has come to have such beliefs (if they do have them), such false beliefs are not themselves illusions (to call them such is a misnomer—illusions are perceptual).

experience to a portion of the foveal region. Alternately, the scope flaw formulation requires that our vision is restricted to the foveal region (or whatever region to which we are currently attending—it is generally, though not necessarily, the foveal region), which is prevented by the structural flaws, which entail that we *can* experience all of the structural flaws at once. Clearly, the result is two contrary formulations. Thus, we cannot maintain both kinds of flaws simultaneously.

3.4 ALTERNATIVES TO THE INCOHERENT FORMULATION OF THE PREDICTED NATURE OF EXPERIENCE

We have a few options. We could reject one of the kinds of flaws, establish that there are actually *two* (or more) distinct illusions, or reject the Grand Illusion hypothesis in the hope that some promising new direction might be found upon its deconstruction. Ultimately, the third option is the only viable one, because there is no principled reason for selecting one kind of flaw and rejecting the other, since they are both taken to be illusory. Furthermore, arguing that there are at least two distinct illusions will also fail because, ultimately, they are illusions of the same perceptual system, such that they will end up with the same predicted nature of perception, whether they are considered as one illusion, or separately.

It is worth being thorough in our inquiry regarding whether we could consider our failure to notice one of the kinds of flaws as an illusory experience, but not the other kind. Since it is only when we consider the together that the formulation of predicted nature of perception becomes incoherent, it seems reasonable to assume that, if we could successfully exclude one of the kinds of flaws, then formulation would no longer be problematic (at least with respect to internal coherence). In order to be able to differentiate the two kinds of flaws in this way, we need to find some theoretical grounds for arguing that the feature in virtue of which one formulation is

illusory is absent from the other formulation. This means that we must evaluate on what grounds we are assuming that each are illusions.

Let us begin with the structural flaws; recall that the expectation which motivates the claim that our failure to experience the flaws is illusory is generated by both our knowledge of the functional structure of the visual system as well as our ability to experience in some of these flaws under contrived circumstances. Thus, the claim that our experience under normal conditions (during which we fail to experience these flaws) is illusory stems from the assumption that, since these flaws are generated by the structure of our visual system, and the functional structure of our visual system is what processes our visual information, then it seems that this “hardware” would determine our experience.

Similarly, the intuition that our failure to notice the scope flaws (that attention restricts what we perceive) indicates that we are subject to an illusion rests on the failure of our visual experience to conform to the apparent constraints of the operations of our visual system. It is argued that it seems to us as though we notice so much more than we do about what is going on before us, and it is in this sense that our experience is said to be illusory.

Assume for the sake of argument that, if these are illusions, they are so for the reasons argued. We must compare the reasons that each are considered to be illusions in an attempt to determine whether there is any way that we could differentiate the two kinds, such that one is an illusion but the other is not. In the structural flaws formulation, we are subject to an illusion because our experience does not conform to what we would expect (given what we know about the structure of the visual system). Likewise, in the scope flaw formulation, we are subject to an illusion because our experience does not conform to our expectations (given what we know about the

attention-dependent nature of perception). Both cases share a common formulation: the perceiver is subject to an illusion because their experience does not conform to what we would expect their experience to be, given the kind of flaw under discussion. Furthermore, to deny that one is an illusion, is to argue that it cannot be an illusion on these grounds (it follows that if it is not an illusion, then it is not an illusion for these reasons, as well as any other reasons). But this would serve as a counterexample to the formulation on which the other illusion depends, such that it would also rule out the other kind as an illusion. As such, we cannot maintain that one is an illusion, but not the other.

Nor can we argue that, rather than both being parts of the Grand Illusion, they are two separate illusions. Suppose that we wish to argue that they are distinct illusions, in an attempt to avoid the problematic formulation of the predicted nature of perception when the two are considered part of the same illusion. The problem with this is that, ultimately, they are illusions of the same perceptual system, such that, if both are illusions, then the same predicted nature of perception is entailed—they problematic formulation would still arise.

It appears that our assumptions will not allow for a coherent formulation of the predicted nature of perception. As such, it appears that we ought to abandon the Grand Illusion hypothesis in search of a new set of assumptions which will allow for a viable formulation of our expectations regarding predicted nature of perception.

4. A NEW FORMULATION OF THE PREDICTED NATURE OF EXPERIENCE

4.1 CONSIDERING THE AFFECT THAT EACH KIND OF FLAW HAS ON OUR EXPERIENCE OF THE OTHER KIND OF FLAW

Suppose for the sake of argument that my deconstructive project is successful—where do we go from here? In other words, though deconstructive projects are useful in their own right, it may be fruitful to search the ruins of the old problem for insight into direction for new position. This appears to be just such a case.

Recall my argument for why we cannot coherently formulate the Grand Illusion hypothesis. I alleged that the Grand Illusion hypothesis resulted from faulty expectations regarding the predicted nature perception based on considering only isolated bits of data. When considering the data from the structural flaws, we assume that since we can experience, say, the blind spot, then our failure to do so under normal circumstances must be due to our being subject to illusion. Likewise, from the research on the scope flaws, we infer that, since the Ordinary Perceiver is surprised when she discovers that she fails to notice large, non-occluded events that are happening immediately before her, it must be the case that it seems to her as though she is receiving visual information from the entirety of her visual field (that is to say, from outside of the restricted area to which she is attending), when in fact this must be illusory, since we know she does not have any such experience of information external to the area to which she is attending through her failure to notice such events. From this, we have generated two contrary sets of expectations regarding what the predicted experience of perception must be like, such that deviation is grounds for arguing that such an experience is illusory. However, as we

saw earlier, this results in an incoherent formulation of predicted nature of perception. When we consider the two individual formulations together, we discover the two cannot be maintained simultaneously.

This failure is insightful, however; by examining what went wrong, we are able to see an alternative formulation of our expectations regarding our experience, given the functional flaws. The most important lesson here is that we must consider the data from the research on both kinds of functional flaws simultaneously in order to more accurately develop our expectations. The attention-dependent nature of perception is equally important of a consideration as the functional structure of visual system. For example, though we receive no visual information from the region which falls within the blind spot, given the attention-dependence of vision, namely we only notice that to which we are attending, which usually (though not necessarily) is that which falls within our foveal region, as well as the fact that are blind spot falls outside of that region, then we should not expect to experience it under normal conditions.⁶² Likewise, though it is true that we fail to notice large scale events that happen immediately before us, this is exactly sort of thing we should expect, given what we know about the attention-dependent nature of perception; if we are not attending to the change, we fail to notice it.

The previous incoherent formulation appears to arise due to emphasis on only a small portion of the causal sequence of perception. These two kinds of flaws are distinct with respect to their causal role in perception. In other words, there is more to consider with respect to what we experience than the means by which we process the visual information; quite literally, we do not “read it off the retinal image”. Instead,

⁶² Though it does seem reasonable to expect that we can experience it under conditions in which we are attending to that region. For example, we are able to experience our blind spot when we have closed one eye and, though we keep our gaze steady, we are attending to some image or symbol within that region.

there are other factors that determine our experience as well.⁶³

We can now begin to see a new formulation of our expectations when based on simultaneous consideration of the affects of both kinds of flaws on perceptual experience. If we note that the attention-dependent nature of perception prevents us from noticing that to which we are *not* attending, and we are (almost) always attending to the foveal region (which we *do* experience as detailed, distinct, in color, etc.), then it follows that we should expect that we would fail to notice the lack of detail, color, etc., outside the foveal region. As such, it appears as though our experience does conform to our expectations, and, thus, we are not subject to a Grand Illusion.

Once we have amended our expectations of what our experience must be like based on simultaneous consideration of both kinds of flaws, we see that our experience no longer fits the old framework for an illusion: our experience no longer fails to conform to what we would expect (given not only our knowledge about the flaws, but considerations of how they impact one another).

4.2 EXPLAINING THE SURPRISE WITHIN THE NEW ACCOUNT

However, I still owe an explanation for why it is that the Ordinary Perceiver is surprised when she experiences her flaws. First, remember that under the new formulation, visual experience conforms to our (the theorists') expectations. It does not follow from that, that the Ordinary Perceivers expectations will be the same as those which we have come to have the virtue of reevaluating the assumptions of the Grand Illusion hypothesis. She will come to form her expectations in a different way, namely through how her experience seems to her. We are now in a position to

⁶³ Note that I am *not* suggesting that there is some ultimate perceptual product—some finished object or image which is then perceived. Rather, I am arguing that there is more involved in determining what is the nature of our perception than the narrow evaluation of the flaws suggests.

understand why, though she is not subject to the Grand Illusion, her visual experience is such that we are nonetheless tempted to call it "photograph-like": perceptual experience under normal conditions is, due to the attention-dependent nature of perception, restricted to the visual information acquired through the foveal region. This region contains a high density of cones, allowing for her experience to be highly detailed and colored. Note that I am not positing any mechanism which somehow prevents her from experiencing the structural flaws; they are unhidden at all times, available to be perceived at any time to which they are attended but, due to the attention-dependent nature of perception, she doesn't notice them under normal circumstances⁶⁴.

If our newly-formed expectations regarding the predicted nature of perception are roughly correct, then her experience is of visual information acquired through the foveal region, which is everywhere detailed, distinct, continuous, and in color. Thus, anytime her vision does not seem that way to her, she is surprised, not because she was expecting the region containing the flaw to conform to a "photograph-like" experience of perception due to the fact that she previously experienced that region as such, but rather because she was not expecting it to fail to conform with her previous visual experiences (which, though not of that region, nonetheless fail to indicate that experience of the flaw regions would be any different). Even if we want to argue that, not only did she not expect that experience, but she also had expectations of what that experience would be like, all we could say is that she expects the same quality of detail and color that she normally experiences.⁶⁵

⁶⁴ During normal perception, she is attending to the foveal region, while the structural flaws which we have been discussing are located outside of the foveal region.

⁶⁵ Given the treatment of filling-in found within at least some aspects of the debate, it is worth mentioning that once our expectations have been reformulated in the proposed way, there is no longer the burden of explaining why we do not experience our flaws during normal perception. This has two potential implications for the advocate of filling-in: (a) some may abandon this sort of explanation since we no longer need to account for our failure to experience the flaws, or (b) some may accept the

4.3 CONCLUDING REMARKS

In conclusion, rather than expecting to experience the flaws under other conditions, our expectations can be modified by considering other, relevant constraints on perception. Attention, for example, is an important constraint: we perceive that to which we attend (and, surprisingly, little else). The structure of the eye is another important constraint, namely regarding the distribution of photoreceptors. How the information is interpreted (how we represent it as being) affects how it seems to us as well. Using these three examples, I will illustrate how constraints can alter our expectations.

We (at least generally) only notice that to which we are attending. This means that, when we are purposefully trying to find our blind-spot,⁶⁶ we are attending to that area of our visual field and noticing features of it in greater detail (or lack thereof) than when we are attending to other regions. However, we are normally attending to the foveal region of our vision. The blind-spot falls outside of that region, and is thus generally outside of our attention. Therefore, the blind-spot need not be altered for us to not perceive it—we may not perceive it because, like the rest of perception, we only notice that to which we are attending. Similar lines of thought have been used in theorizing about change-blindness and inattentional-blindness, and I see no reason why we cannot apply this explanation more generally.

The distribution of photoreceptors is another important constraint—the more photoreceptors, the greater the visual detail. Since there is a greater number of

new formulation, knowing that they no longer bear the strange explanatory burden of why we are able to experience our flaws under non-normal circumstance. Either way, it seems that, in accepting the new formulation, filling-in has been vindicated from what was generally taken (perhaps erroneously, since there appears to be no discussion of such a theory in its own right, but only as a means to this particular explanatory end) its previous role in the debate—accounting for why we don't experience the flaws—which lead some members of the debate to interpret filling-in as an unreasonable position.

⁶⁶ I do not mean for this to be the only case—there are other ways of discovering your blind-spot.

photoreceptors (and a great deal of those being cones) in the foveal region, which falls in the center of the visual field, then whatever falls into that region of the visual field will be experienced in color and great detail. We generally attend to that which falls within our foveal region, so everything to which we attend is generally experienced in color and great detail. Furthermore, the density of photoreceptors (especially cones) becomes less moving in the direction from the foveal region to the periphery. This lessens the detail and color of our vision. However, we do not generally notice this lessening of detail and color *because we are not attending to that region*. If we normally attend to the center of our visual field, and the center of our visual field is experienced as being in color and having good detail, we (a) do not notice that it is otherwise elsewhere in our visual field and (b) that which we are perceiving *is* experienced as detailed and in color (this is not illusory—it appears reasonable to argue that one’s experience really is that way in this region), which may lead us to infer that our vision (perhaps entirely) is that way. This inference is not subject to the illusion, because it is not the case that we are experiencing our *vision* as everywhere detailed, distinct, continuous, in color, photograph-like, etc. Rather, we experience the region of our visual field to which we normally attend (and which *is* that way) *as* being that way, and could mistakenly infer that our entire visual field is that way (since we do not generally attend to that which falls outside of it—we just turn our heads to look if we wish to see something that falls outside of the foveal region—and, due to photoreceptor distribution, which yields poor information outside of that region, we do not experience anything to disconfirm our experience until we experience the flaws). Given these modified expectations, it seems clear that we are *not* subject to the Grand Illusion.

BIBLIOGRAPHY

- Cohen, Jonathan. 'The Grand Grand Illusion Illusion.' *Journal of Consciousness Studies*. (9:5/6) pp. 141. 2002.
- Dennett, Daniel. 'How Could I Be Wrong? How Wrong Could I Be?' *Journal of Consciousness Studies*. (9:5/6) pp. 13. 2002.
- Dennett, Daniel. 'Filling In Versus Finding Out: A Ubiquitous Confusion in Cognitive Science.' <http://ase.tufts.edu/cogstud/papers/fillin.htm>, last accessed August 12, 2004.
- Dennett, Daniel. 'Seeing is Believing—or is it?' in K. Akins, ed., *Perception Vancouver Studies in Cognitive Science*, vol. 5: Oxford University Press, 1996 pp. 158-172.
- Noë, Alva. 'Is the Visual World a Grand Illusion?' *Journal of Consciousness Studies*. (9:5/6) pp. 1-12. 2002.
- Pylyshyn, Zenon. 'Mental Imagery: In Search of a Fundamental Theory.' <http://rucss.rutgers.edu/faculty/zpimagery.htm> Last accessed August 12, 2004.