Title
Looking and Learning: Pictorial Representation and Visual Skill

Author
Dolan, Caitlin

Publication Date
2020

Peer reviewed|Thesis/dissertation

Looking and Learning:
Pictorial Representation and Visual Skill

By

Caitlin E. Dolan

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Philosophy

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Hannah Ginsborg, co-chair
Professor Alva Noë, co-chair
Professor Mary Ann Smart

Fall 2020

ABSTRACT

Looking and Learning: Pictorial Representation and Visual Skill

Seeing can be difficult. This dissertation aims to bring out the philosophical significance of this commonplace fact. By examining the roles of difficulty, agency, and skill in visual experience, it sheds light on two related sets of philosophical questions: one about our ability to perceive the world, the other about our ways of representing it.

The first two chapters get at the basic structure of the experience of seeing. I approach this by concentrating on the notion of an opportunity to see. This notion is a familiar one, but upon examination it makes trouble for a traditional and tenacious idea, that seeing is the automatic upshot of a certain mechanical process – a causal chain linking things in our environment to our perceptual organs and our minds. Having presented this conflict, I examine the alternative framework for characterizing visual experiences that the notion of opportunity invites. I agree with other discussions that opportunities to see are defined by contextual conditions like lighting and the placement of occluding objects, and the degree of salience things have in virtue of the kinds of objects they are. But I argue that in order to understand the difference between an opportunity to see and an occasion of seeing, we must recognize skill on the part of the viewer as an additional explanatory factor.

The second chapter turns to a puzzle about how it is possible to see the intrinsic features of things in contrast to the way that they look in certain circumstances – e.g. the color of a painted wall, as opposed to the way shadows darken parts of its surface, or the volume of a cardboard box, as opposed to the square shape of its facing side. This puzzle construes lighting and opacity as *obstacles* to seeing, and my response emphasizes that illumination is an *enabling* condition for vision, and a thing's opacity (the fact that one cannot see past its surface) is what *makes* it visible in the first place. At the heart of the puzzle is the assumption that whenever we see, we see how things look, and the idea that seeing how things are shaped and colored must involve more sophisticated skill than seeing how they look in the conditions in which we encounter them – if it is possible at all. I argue that this gets it backwards. Seeing how things look – e.g. how the play of light affects the appearance of a thing's color, and how different parts of a thing's surface appear from different vantage points – is the more sophisticated business.

The rest of the dissertation focuses on the art of pictorial representation, as a practice that exemplifies the skill involved in discerning the way things look. The third chapter takes up Richard Wollheim's influential thought that seeing a picture is an experience whose nature is twofold, an experience of both a marked surface and a depicted object. I pinpoint the insight in this idea by explaining the duality in terms of the distinction between a marked surface and the way it looks. This allows us to avoid thinking of the experience of the depicted object as an illusion, or as a "vision" in an occult sense, of something in a mysteriously distinct realm. It also brings out the fact that pictures are things whose visual appearances are designed to be especially salient, and are products of an intention to *show* how something looks.

The fourth chapter turns to the relationship between depiction and visual resemblance. A central motivation for thinking that the concept of depiction in some way implicates that of visual resemblance is an interest in capturing the difference between

pictures and linguistic expression, as two forms of representation. Nelson Goodman criticizes this way of comparing them, in part by arguing that it relies on a flawed conception of the experience of seeing, according to which it always involves an awareness of a determinate configuration of patches of color. I examine attempts to refine the resemblance theory of depiction to avoid Goodman's objections, and argue that they too build in problematic assumptions about visual experience. They require that depiction always presents viewers with forms of appearance that they are already familiar with – aspects of shape and color that they already associate with things of a certain kind. But we should not think of depiction as constrained in this way. The role of shape, line, and color in depiction is not to remind viewers of looks they've seen before, but to get them to see, possibly for the first time, how things can look.

The final chapter attends to the activity of drawing from life, or creating a depiction based on a model. This activity serves as one final illustration of the idea that visual skill is developed with effort. The first part of the chapter argues that we should not think of the process of drawing from life as constituted by two stages, first seeing how the model looks, and then manually modifying a drawing surface to convey that look. Rather, the drawing process involves both visual and manual skills throughout, and sometimes, it is only through the process of drawing that the artist comes to see how the model looks. The second part of the chapter clarifies the nature of the skill exhibited by drawing from life by examining how artists typically train to do it. Finally, I present a way of thinking about pictorial realism in terms of the kind of visual skill that drawing from life requires. On this way of thinking, realistic depiction is not a matter of capturing some particular aspect of the visual world, but rather a matter of conveying our visual sensitivity to the contingent features of objects in our surroundings, whatever they may be.

Preface

The idea that unites the chapters of this dissertation is that seeing can be difficult. This commonplace fact is often neglected by philosophers, but it is crucial to understanding many philosophical questions about visual experience and visual art. Each of the five chapters that follow show how philosophical puzzlement on a particular topic is illuminated by recognizing the fact that seeing takes skill. By attending to the relevance, for a variety of phenomena, of the idea that seeing is predicated on skill, I aim to engender philosophical understanding of the fact that visual skill can increase in degree, can develop in various ways, and can be exhibited in the things we make to be looked at.

The first two chapters focus on visual perception as a general means of navigating the world. In approaching this topic, I have found it useful to foreground the phenomenon of opacity, making clear how it is fundamental to the structure of the visible world, and I articulate basic aspects of the visibility of the world by drawing connections between opacity and shape, color, and illumination. This provides a conception of sight as a capacity to realize the potential the world has in virtue of being visible. In the first chapter, I focus on the relationship between opportunities to see and the event of actually seeing something, with an eye to clarifying the general role that the notion of skill plays in our thought about visual perception. In the second chapter, through reflection on the relationship between opacity and the visual appearance of shape, I arrive at an idea which is central to what follows: that there is a basic distinction between the skill it takes to see something, and the skill it takes to see how a given thing looks. I argue that the latter is more sophisticated than the former, by explicating the way a thing looks as a matter of how its features structure the range of opportunities there are to see it. This distinction, and the special skill involved in seeing how things look, becomes my primary tool for making vivid the fact that seeing can be difficult, and some objects of sight are harder to see than others.

A different approach would be to focus on how the development of visual skill enables us to see further and further kinds of things or features of things, rather than the skill of seeing how things look. One could point to the keen eye of a geologist who can spot differences in the mineral composition of the rocks that her hiking companions see as homogenous, or the expertise of a rowing coach who can see that the angle of a stroke is slightly off. I have not pursued this here, in part because not everything can be pursued in every dissertation, but also because things' looks have an especially interesting status in relation to vision. For one thing, the possibility of seeing how things look (whatever they may be, or look like) is entailed by the possibility of seeing anything at all. We can only see rocks and rowing strokes if such things exist. But whatever there is for us to see, it is possible to see how it looks. Looks are also the subject of a tempting conception of what it is to see: that in seeing, we are immediately made aware of things' visual appearances. This way of thinking recognizes that the way things appear is distinct from the way things are, and sometimes conceives of visual appearance as a two-dimensional array of colors and shapes. Though this way of thinking is tempting, it is also uncomfortable, since it can seem to raise a difficult question about how this awareness relates or contributes to our awareness of visible things themselves. I am interested in what is so uncomfortable about answering this question, and how the context of the tempting conception of seeing affects the way it is posed. What I say in chapter 2 suggests that we can understand this better by examining how, exactly, the capacity to see entails the possibility of seeing how things look.

These questions are about both the epistemology and aesthetics of perception, and

are tied to certain topics in the philosophy of art. Visual art, especially pictorial art, is often characterized as undertaking an investigation into visual appearance or the visual world, but interpreting that claim requires saying something about what the visual world is and how it can be investigated. I think that the topic of visual art provides an excellent lens for these philosophical questions about seeing, and perception in general. Thus the last three chapters focus on the topic of depiction, or pictorial representation, and examine both what it takes to make pictures and the kind of visual experience that they afford. It is not hard to demonstrate that making a picture, at least making one by hand, requires skill. This will not seem controversial to most people, insofar as we associate depiction with artistic talent. We are rightly impressed with the rich history of pictorial art, but also, the education many of us receive barely includes any opportunities to practice and improve our own picture-making. The idea of visual skill most often becomes salient in the context of thinking about painters, sculptors, designers, and others who make things that we value for their visual appearance. But though the ability to make pictures may be a "gift" in some sense, it is philosophically important to recognize the fact that it involves skills that are developed through certain kinds of activities and encounters – and that pictures themselves afford this kind of development for those who look at them. Chapters 3-5 aim to bring this out.

Chapters 3 and 4 address, respectively, two very influential ideas in the philosophy of depiction: one, that the visual experience afforded by a picture is marked by a distinctive duality of pictorial surface and depicted object, and the other, that pictures must resemble what they depict. Chapter 3 connects the ability to see something as a picture with the phenomenon of seeing how things look. Chapter 4 argues that pictures do not just impress the viewer with the creativity and technical ability of others, they also expand the viewer's own visual ability, by *showing* them – getting them to see – how something looks. These chapters approach depiction as philosophers tend to, from the perspective of the beholder of a picture, rather than the maker of one. The final chapter, however, centers on the perspective of the picture-maker, and attempts to clarify the skill involved in making a picture that is based on an object or scene one sees. In a way, the most vivid appreciation of this skill comes from looking at pictures themselves. But I hope to have brought out the philosophical significance of this phenomenon, in the abstract.

In writing this dissertation, I have benefitted enormously from the help and inspiration of many people in the Philosophy Department at Berkeley and beyond it. First and foremost, I am immensely grateful to my advisors, who have provided invaluable guidance throughout this long and winding process. Hannah Ginsborg, Alva Noë, and Barry Stroud have shaped my understanding of what philosophy is, and I cannot imagine what it would be like to do philosophy without having known them. They have always encouraged me to write in my own voice about the questions I genuinely find interesting, and the way they each pursue philosophy in their own style has shown me what it means to contribute to this tradition. I am indebted to their work in ways that I have only begun to comprehend. It has been an honor to be their student, and also a great pleasure.

My thinking on the topics of these chapters has been developed by the many connections I have made as a graduate student at Berkeley. In particular, I owe thanks to Mike Arsenault, Adam Bradley, Peter Epstein, Alex Kerr, Umrao Sethi, and Klaus Strelau, for participating in a reading group on perspective that greatly improved my understanding of the geometry of optics and its relevance to visual phenomenology, and for stimulating conversations in many other contexts. I am also grateful to the members of the Townsend Working Group on Aesthetics, including Joe Kassman-Tod, Dave Suarez, B. Rousse,

Samantha Matherne, Janet Broughton, Tyler Haddow, and Tim Clarke, for all of the insights into pictorial style and photography that they shared in our meetings. The members of Alva's working group provided helpful comments and discussion on several chapter drafts. Jim Hutchinson, Kirsten Pickering, and Yuan Wu have also been extremely generous interlocutors and companions in art appreciation. My work has been enriched by the content of all of these conversations, and I have been enriched by the sense they created of a community with a shared endeavor.

I would also like to acknowledge the formative and productive conversations I have had with Maarten Steenhagen and with Jenny Judge. I am grateful to both of them for their enthusiasm for questioning even the most entrenched presuppositions of a philosophical topic, and for their patience and creativity in talking through things together, from the ground up. I am also extremely grateful to Chase Booker, who has introduced me to so much art, and lent me a copy of *Art and Illusion* that I have been rereading ever since.

Finally, I would like to thank Luke Jensen, for all the light he sheds on whatever he turns his attention to, and my family. I could not have done this work without their support, encouragement, and understanding.

Contents

Chapter 1
The Opportunity to See: Skill and Visual Salience

We are inclined to think of the phenomenon of sensory perception as part of an account of how we come to know about the way things are. We take it, for instance, that we find out about what the world is like when we look and see. This way of accounting for our status as informed creatures suggests that the world contains two kinds of potential. On the one hand, we see because the world is visible – because *it can be seen*. On the other, we see because we have the power of sight – we see because *we can*.

Both of these kinds of potential are realized when we see, whatever and whenever we do. In that sense, seeing can be thought of as an interaction between us and the things around us, as something that happens when a sighted individual "meets with" some of the visible things in the world. It is tempting, or at least traditional, to apply a certain model of causal explanation to this interaction. On this way of thinking, episodes of seeing are instances of causal chains of a certain kind, beginning with the things in the world that we can come to see, and ending with a psychological state induced in an individual who thereby counts as a viewer. The visibility of things is then a matter of a kind of causal efficacy they possess, a power they have to impinge on bodies with physiological constitutions like ours, triggering a process that culminates in a state in which a person *seems to see something*. Visual phenomenology, or any analysis of the *experience* of seeing something, is a matter of grasping what happens to us at the final link in this causal chain.

From one perspective, the idea that seeing is a causal interaction between an individual and the world can seem to be required by the conviction that it is a way of becoming informed, but from another it can seem incompatible with it. John McDowell, for one, has argued that we can't accommodate the epistemic role of perception unless we conceive of it as drawing on our active capacities for reasoning (McDowell 1994). This might be taken as an expression of the thought that when we say someone *can see*, we are not referring to any tendency she has which merely makes certain things likely to happen to her, but are rather crediting her with an *ability*. But, like the thought that seeing is a causal interaction, this idea can seem to create as much trouble for the epistemic significance of seeing as it resolves. If seeing is a manifestation of ability, then it would seem to be a way of exerting an influence on the world, rather than a way of appreciating how it is already.

So what are we to think? Is seeing something we do, or something that happens to us? How should we deal with the fact that it seems to be both? The aim of this chapter is to recommend that we hold on to the conviction that it *is* both, and to provide a way of doing so. According to the conception I will explicate, visibility is a matter of affording a certain kind of opportunity, and being sighted is a matter of being able to exploit opportunities of that kind. I will make room for the idea of exploiting an opportunity to see, as a way to become informed about how things are objectively, rather than an occasion for influencing the way the world is. Seeing isn't a choice, but it is nonetheless the exercise of a certain kind of ability.

The structure of this chapter is as follows: in section 1, I make some general observations about the concepts of ability and opportunity, sketching the explanatory framework that gives them significance. In section 2, I present the debate over whether to apply this framework or the causal one to the case of seeing, by way of John Hyman's work to defend the former. In section 3 I identify a problem with the way Hyman elaborates this framework by invoking the notion of visual salience, and in section 4 I try to resolve it by

clarifying the relation between salience and visual skill. Getting clear on how to understand seeing as the realization of ability and opportunity involves recognizing two main points. The first is that seeing is the manifestation of skill in meeting a challenge of a particular kind. Sometimes it is relatively easy, but it is the sort of thing that can also be difficult. The second is that the visibility of the world, as its affording opportunity of a certain kind, is shaped by an indefinitely broad range of features.

1. The concepts of ability and opportunity

Given the aim of clarifying the kind of potential that we could be talking about in saying that someone *can see*, and that something *can be seen*, it is helpful to start out with some general points about potential and its varieties.

There are many different kinds of possibility or potential that our use of the word "can" is able to convey.[1] Among them are logical possibility ("Modus ponens cannot lead from true premises to a false conclusion") and legal possibility ("Americans can vote once they turn 18").[2] These are distinct from the kinds of potential at issue in this paper. My interest here is in the "can" of ability and the "can" of opportunity. Examples of the former include: "Martha can speak Spanish"; "Stephanie can drive a stick-shift." Here are some examples of the latter: "In January when the lake freezes over, you can skate on it"; "One can buy a decent cup of coffee at the café around the corner."

We can make progress understanding the nature of ability and opportunity by considering the relations they bear to each other. Abilities can only be exercised when the opportunity arises (or "presents itself," as Hyman puts it); opportunities can be realized only by those with the appropriate abilities. This talk of abilities "appropriate" to opportunities marks an intimate connection between the two concepts. Any opportunity is an opportunity to exercise abilities in certain ways, and the possession of an ability entails the possibility certain kinds of opportunities arising for its possessor. Abilities and opportunities can thus be individuated or classified in terms of each other: abilities in terms of their "opportunity conditions," and opportunities in terms of what it takes in the way of ability for them to be realized.[3]

There are two important things to note about these ways in which ability and opportunity are related – about the fact that abilities can only be exercised when the opportunity arises, and that opportunities can be realized only by those with the appropriate abilities. The first is that they themselves employ the word "can." This is neither the "can" of ability nor opportunity, but "can" in a sense which expresses a further kind of potential. Hyman calls it the "all-in 'can'," and Kenny refers to as an "overall sense of 'can'." The conjunction of the possession of an ability and the encounter with an opportunity contrasts

[1] This discussion of possibility and ability draws on those found in Hyman 1994 and Kenny 1992. It also takes inspiration from Noë's discussion of presence (Noë 2012). Ryle's discussion of know-how is in the background of much of this thought about cognitive and perceptual ability (Ryle 1949).

[2] One might wonder whether all of these classifications of possibility are on a par, whether some are more fundamental than others, and whether the less fundamental can be understood in terms of the more fundamental. I will leave unaddressed many questions about the way different kinds of possibility relate to each other. But what I will say entails that the "cans" of ability and opportunity both divide into different species. I will show that there is a difference between the sort of potential expressed by "John can swim" and "John can see," though they both refer to abilities. I will also show that there is a corresponding difference in the kind of opportunities expressed by "With a wetsuit, you can swim in the bay in February," and "On a clear day, you can see the marina from this window."

[3] Together, opportunities and abilities constitute the enabling conditions for occasions of certain sorts – such as swimmings, skatings, and seeings.

with "thinner" kinds of potential, when there is opportunity without ability, or ability without opportunity: "The island can be reached by boat, but Thomas will never get there since he barely knows how to sail"; "Markus can make an excellent coffee cake, but he won't today because there's no butter in the house."

The second thing to note is that occasions on which these kinds of potential can be realized are described as both the exercise of abilities, and the arising of opportunity. We can understand the potential represented by abilities and opportunities by considering the nature of the episodes that count as their realization, but we need to be careful to distinguish two kinds. On the one hand, there is the arising of opportunity for a person on an occasion – a circumstance described using the overall "can." As I've explained, we can characterize kinds of abilities and opportunities by recognizing their relations to each other, how their combination constitutes the distinct potential expressed by the "all-in 'can'". Though this circumstance realizes two individual potentials in a certain sense, the *exercise* of an ability counts, in a sense, as a fuller realization of them. We can understand what is expressed by the overall sense of "can," as well as the ability and opportunity it implicates, by considering the nature of the episodes that count as exercises, or as the full realization of these kinds of potential.

This brings us to a further point, which will play a pivotal role in the debate over whether and how to apply these notions to seeing. When we say that someone *can X*, in this overall sense of "can" – when we mean that she has the ability to X, and has encountered the opportunity to X – we are not yet saying that she in fact *does* X. Often, it is not only the case that someone *can* X, but that Xing actually occurs. But the statement that someone can X (in this sense) does not refer to an occasion on which she does. It refers only to a *meeting* with opportunity, not to what comes of it.

This is important because it provides a way of differentiating the exercise of ability from the manifestation of another kind of potential, a kind which Hyman refers to as a "natural power" or "tendency." This kind of potential is *necessarily* realized when preconditions are met: a flame can burn wood, insofar as when it is placed in close enough proximity with sufficiently dry wood, the wood burns. The satisfaction of the relevant preconditions entails that the potential in question *is* realized or manifested, not simply that it *can* be (in some further sense of "can"). In contrast, the "cans" of ability and opportunity do not relate to the manifestation of these kinds of potential in that way.

Before moving on to address the case of seeing, I want to put forth one more way of getting a handle on the nature of a particular opportunity or ability, namely by way of the factors that can enhance and interfere with it. For example: ComiCon is rife with opportunity to talk about *Star Trek*, the Eastern APA less so. My ability to follow French films without subtitles is diminished when I'm sleepy.

Hyman touches on this in discussing the basis for distinguishing the roles played by ability and opportunity in an episode that counts as their joint realization. He claims that while ability is always a "positive explanatory factor" for an occurrence, the notion of opportunity is typically "no more than a negative factor, the absence of circumstances that would prevent or interfere" (Hyman 1994, 240). But it is important to note that this is a contrast in how the explanations of certain occurrences tend to be *characterized*. All it means is that typically, we refer to opportunities in a negative way, while we point out abilities positively. However they are picked out, though, both play a positive explanatory role in a certain sense – their presence or possession explains what does, in fact, happen. Moreover, it seems possible to characterize both abilities and opportunities either positively or negatively, even if one of these options is more typical. We can say that on an occasion someone could

reach the island by boat because they left well before the storm arrived and didn't get caught in it (a negative characterization of his opportunity), or because the wind was in their favor (a positive one). We can say that a gymnast made a comeback in the final round because they were so skilled at the balance beam (a positive characterization of her ability), or because they didn't let disappointment at the beginning of the competition distract them later on (a negative one).

Finding general principles for distinguishing the roles of ability and opportunity in the events that realize them proves extremely difficult, at least at a high level of abstraction. Hyman reports the suggestion that abilities are "internal" while opportunities are "external," admitting that this metaphor is very hard to explicate further. He leaves it with the suggestion that "an opportunity has to do with an agent's situation or circumstances, rather than his condition" (Hyman 1994, 240). But it is not easy to articulate what makes us count something as affecting someone's "condition" and incapacitating her, rather than interfering with her "circumstances" and robbing her of opportunity. It seems we can only get so far in trying to give general criteria for the distinction. However, it does seem to apply clearly enough in many concrete cases. Intuitively, the gymnast's anxiety affects their ability, while the weather affects the sailor's opportunity. So, with the intuitive difference on the table, we are prepared to delve into a somewhat more specific class of occurrences: episodes of seeing.

2. Accounting for seeing

As mentioned in the introduction, there is a long philosophical tradition of conceiving of the occurrence of seeing as the instantiation of a certain kind of chain of events, implicating both would-be viewers and the would-be visible things around them.[4] In "Vision and Power," Hyman takes issue with this picture by making use of the above distinctions between types of potential and the sorts of explanatory frameworks that they participate in. According to him, the causal theory articulated by Grice and accepted by many others is mistaken because it "misrepresents its *analysandum* as a natural power" (Hyman 1994, 236). Hyman proceeds by trying to show that the phenomenon of perception in fact fits into the explanatory framework of ability and opportunity, rather than that of natural power or tendency.

Though the target of Hyman's objection goes by the title of the "causal theory of perception," this debate need not be construed as a disagreement about whether perceptual episodes admit of causal explanation. This objection to the "causal theory" entails that perception is not a causal phenomenon only if all causation unfolds on the model of the manifestation of a natural power. To repeat, the key to the difference between the manifestation of a natural power and the exercise of an ability is found in the fact that the former is *necessitated* when preconditions are met, whereas the latter is only *possible* (in a certain sense) when the relevant ability and opportunity are present. This is what sets the terms of the debate between Hyman and the proponent of the causal theory. Their common ground is a general sense of the kind of factors that make a difference to whether and what people see. The disagreement is over the kind of difference they make: whether they structure opportunities and qualify abilities, or instead determine whether preconditions are met. Whether or not the idea of causation is broad enough to include both of these

[4] This chapter focuses on Grice, but other articulations of a causal theory include Strawson 1974, Strawson 1979, and Noë 2004. Further criticism of this kind of theory includes Snowdon 2009.

explanatory structures is a further question, which I will not address here.[5] For the sake of brevity, however, I will continue to use "causal connection" and related terms to indicate to the specific relationship exhibited by a natural power and its manifestation.

Grice's arguments for his particular causal analysis of seeing provide a starting point for adjudicating this issue. They take the form of certain counterfactual claims about seeing and not seeing, which Grice takes to show that it is "a necessary and sufficient condition of its being the case that *X* perceives *M* that *X*'s sense-impression should be causally dependent on some state of affairs involving *M*" (Grice, quoted in Hyman 1994, 237). *X*'s sense-impression is something that happens to *X*, namely its seeming to *X* "just as if he were seeing something of such and such a kind" (Hyman 1994, 237). The cases are the following:

> *Hypnotism.* An "expert," using some kind of "apparatus" or "technique," hypnotizes a subject so that it will look to him as if there is a clock on a particular shelf, even when there is no clock there. If, as it happens, there *is* a clock there, it would not be correct to say that the subject sees this clock when he looks towards the shelf.
>
> *Mirror.* A mirror is placed in front of a pillar so that a numerically different but qualitatively similar pillar is reflected in it. A subject stands in between the two pillars. Though he is facing the first, he does not see the first pillar, but rather the second (the one behind him).
>
> *Hand.* A subject is in good lighting, holds his hand up in front of him, and looks at it. He sees his hand.

The first two are supposed to help establish that causal dependency is a *necessary* condition. Hyman summarizes their lesson as follows: "a mere correspondence between what a person takes himself to see and what is there before his eyes does not establish that he sees what he takes himself to see" (Hyman 1994, 237). Grice uses the third to illustrate that a causal connection that is conceptually sufficient for seeing to occur. Not just any kind of causal relation will do, but the relevant one is to be picked out by reference to this paradigm case.

To the argument for the necessity claim, Hyman responds that "these stories do not establish what they are meant to" (Hyman 1994, 238). It is true, and they do show, that "a mere correspondence between what a person takes himself to see and what is there before his eyes does not establish that he sees what he takes himself to see" (Hyman 1994, 237). But they show this by illustrating that it is possible, in various ways, to prevent someone from seeing something, and also that when such an attempt succeeds, the victim may nonetheless *take himself* to see the thing in question (or something fitting its description). They do not show that what is needed, rather than mere correspondence, is a causal connection.

If these stories fail to establish the causal theory's necessity claim, they also provide a basis for showing that it is false. This is because these cases might well be described as examples of denying someone the opportunity to see something, and interfering with someone's ability to see something. If a case can be made that this is what the stories illustrate, then we will have reason to accept that seeing fits the explanatory framework just described, as opposed to the kind of causal picture that Grice proposes. But the difficult question we now face is how to adjudicate between the two proposals. The cases on the

[5] For a discussion of explanation that points in the direction of the verdict that both of these structures can be used to capture causation, see Black 1962.

table can be neutrally presented as examples of preventing someone from seeing something. What resources do we have for determining which of the explanatory frameworks fits them?

To settle this question, Hyman imagines a debate between "causalist" and a "skeptic." They agree on the following requirement for positing a causal relation:

> *E1* causes *E2* if and only if it is logically possible for *E1* to occur without *E2* occurring and there is an empirical generalization to the effect that whenever an event occurs of a type to which *E1* belongs an event occurs of a type to which *E2* belongs. (Hyman 1994, 245)

Hyman adds that "the generalization must be stated in such a way as to allow for the possibility of interference." This, I take it, is what makes the causal relation to be posited compatible with the logical possibility of *E1* occurring without *E2* occurring. Interference is conceived as a matter of the satisfaction of preconditions: *E1* will occur without *E2* occurring when preconditions are unmet; when such preconditions are met, *E1* causally necessitates the occurrence of *E2*. Relying on our common sense of the sorts of factors that figure in what will *either* count as the preconditions for seeing *or* the opportunities that vision implies, we can then try to assess what should be said when *E1* is the presence of a visible object and *E2* is the event of the sighted individual in its presence seeming, or "taking herself," to see it. That is, we can ask whether it makes sense to make this kind of empirical generalization about seeing things.[6]

In the debate Hyman imagines, the causalist makes a series of claims about what needs to be admitted in an account of the hypothetical fact that someone does not seem to see a visible object in her presence. Many of them point to various features of some hypothetical set of circumstances – states of distraction, darkness, distance – as examples of causal interference with the object's ability to have the relevant effect on the viewer. Others describe the circumstances as those in which the viewer *does* see the thing in question, but doesn't *notice* that she does. The skeptic responds to each by simply pointing out that these claims amount to *insistences* on the causal analysis, rather than defenses of it. For the scenarios of the first kind, we haven't been presented with any reason to think that if these problematic aspects of the circumstances were removed, the viewer would *necessarily* see the thing, rather than simply being *able* to in the "overall" sense. And when it comes to the second, while it might be the case that people *sometimes* see things without noticing that they do, this does not entail that they *must* see the visible things in the circumstances described, as a matter of necessity.

Hyman's presentation of these points is subtle and convincing. But it does not quite accomplish all we might hope to do in adjudicating this dispute. These points may convince us that we haven't been presented with an argument *for* the causal analysis, and so if we didn't feel a need to accept it at the outset, we shouldn't do so in light of anything that has been said. But one might agree that we needn't accept Grice's picture for the reasons provided, but still feel that we are in the dark about why we *can't*. Moreover, Hyman's points don't provide an argument in favor of the alternative either. One might grant that "having

[6] We should not take it that Hyman's "skepticism" involves claiming that *all* kinds of empirical generalization about seeing is illegitimate. It is important to recognize that on an analysis of seeing in terms of opportunity and ability, it is perfectly reasonable to form *expectations* about what people see, on the basis of generalizations about when people are more or less likely to see things. The Gricean causal analysis of seeing goes beyond this, advocating a particular view of what these expectations are based on.

the opportunity and the ability to ϕ do not suffice to guarantee that one ϕs: one needs to take the opportunity and exercise the ability" (Hyman 1994, 248). And one might also grant that if the concept of an opportunity to see is not "otiose," the causal analysis must be mistaken. But one might object that the antecedent of this last conditional has not been proven, and that the burden is on the causalist's opponent to justify it.

Given that the burden of proof may not be balanced in favor of the opportunity-and-ability analysis at the outset, one admission that Hyman makes puts his position on very unstable ground. To appreciate the danger, we must note that if the proponent of the causal analysis is someone who views it as prima facie more sensible than the competing analysis, then in effect he is someone who claims not to be able to understand "can see" in the "all-in" sense, as distinct from both the visibility of an object, and the event of seeing it. The dangerous admission is that "we often use the phrase '*X* can hear (or see)…' to mean '*X* does hear (or see)…'." (Hyman 1994, 248). But *why* do we use "can" to talk about actual occurrences in this way, if not because "can" and "does" simply mean the same thing in this context – if not because there isn't, after all, any conceptual difference between the circumstances being appropriate for something to be seen by someone, and this seeing actually taking place?

To be fully responsive to this position, we need to identify why someone might have trouble with the notion of an opportunity to see – why the suggestion that perception is the manifestation of ability might seem problematic. If it can be shown that whatever generates this worry involves a misinterpretation of the proposed framework, then we will have something to say to this skeptic of the alternative analysis. This won't amount to a deductive argument in favor of this conception of seeing, but it does involve positive work: it requires us to clarify further what the framework really comes to, to get a better sense of what more we might be saying when we say that someone sees something, of how it goes beyond stating that the overall conditions are right for it.

This is the work I will do in the final two sections, ultimately arguing for a substantial revision of Hyman's proposed understanding of the fact that someone sees something on a given occasion. I show that in applying the notions of ability and opportunity to seeing, we must recognize a discrepancy between manifestations of perceptual versus non-perceptual abilities, and so must look for a significantly different way for the idea of ability to get a grip. I argue that what Hyman says on the topic of seeing falls short of this. In the final section, I explain how to do better.

3. Opportunities and opting

In order to properly understand what is required if we take the notions of opportunity and ability to apply the phenomenon of seeing, we need not only to contrast seeing with the manifestation of natural tendencies, but also to consider its relationship to the kind of power or potential that is realized in voluntary choice.

The examples used in section 1 to bring out the notion of opportunity involved potentials of this kind. When it comes to abilities like sailing, driving a stick-shift, and the like, their possession seems to put us in a position to make certain choices. To pick another example: being able to swim entails being able to do something voluntarily, being able to choose to swim when the opportunity arises. Being in a position to choose to swim implies being in a position to refrain from swimming. This is why the ability to swim and others like it earn the title of a "two-way power": there are two ways for this kind of power to manifest, namely, doing something by choice, and not doing something by choice. That is, the choice

that someone faces when they find themselves in the circumstance of its overall possibility should not be construed as a choice between demonstrating their ability, or not. Rather, it is a choice between taking the opportunity and forgoing it: swimming, or not swimming. Either one presupposes the ability, and in that sense demonstrates its possession.

Abstention provides a way to appreciate the difference between having an opportunity and the corresponding ability, and what actually occurs in that circumstance. It makes clear that what happens is, in a certain sense, *optional*, by allowing us to distinguish the options. That is, the phenomenon of abstention offers a route to the concept of an opportunity. But if visibility is a matter of opportunity – if there is such a thing as an opportunity to see – then this can't be the only route. As Hyman puts it, if seeing implies the opportunity to do so, then "the concept of an opportunity is not coincident with the concept of choice" (Hyman 1994, 241). If we construe seeing as the manifestation of the kind of ability that two-way powers represent, then we have to accept that we see things because we choose to. This runs counter to the conviction that it is not "up to us" what we see: we don't choose what we see, in the way we choose what we do when we decide whether to jump in the lake or stay on the shore. But this leaves us with the question of what *is* involved in realizing the potential that constitutes an opportunity to see – if not the *taking* of that opportunity, in a sense that contrasts with voluntarily passing it up instead. That is, we want to know the nature of the difference between having the opportunity and ability to see something and actually seeing it, and having the opportunity and ability and failing to see it – given that it isn't, as in cases like swimming voluntarily, a matter of one's choice.

Hyman approaches this question by asking about what "bridges the gap" between "the ability to see *M* (the 'all-in' can of perception) and *X*'s actually perceiving *M*." His answer is twofold: this gap "can be bridged either by salience, in which case *M* causes *X* to see it, or by *X* noticing what he might have failed to notice – perhaps but not necessarily, as a result of a voluntary action of *X*'s (looking or straining to see; listening or straining to hear)" (Hyman 1994, 246). He dubs perception a "quasi-natural power" in virtue of admitting of this kind of analysis. I think, however, that this suggestion can't be quite right: it threatens to undo any progress that might be made by the suggestion that visibility is a matter of opportunity. In the rest of this section, I'll explain why.

When Hyman says that the salience of an object "bridges the gap" by causing someone to see it, this amounts to the claim that some occurrences of seeing come about in just the way that exercises of natural powers do: as he puts it, as a matter of "natural necessity" (Hyman 1994, 246). In the other case, of insufficiently salient things, Hyman's account seems too close to a restatement of what we want to understand to be much help. In these cases, instead of salience, it is a matter of "noticing." But Hyman says little about what noticing is, and one worries that in this context it can only mean "see." In a footnote Hyman says that "in the visual case, '*X* noticed *M*' entails that *X* saw *M*, and generally carries the implication that *X* did not see *M* as soon as he had the opportunity to do so, or else that *X*, or someone other than *X*, might well have had the opportunity to see *M* without seeing *M*" (Hyman 1994, 247). But if that's right, then all we have learned is that sometimes a sighted person sees something eventually instead of immediately, or that he sees what he might not have: he sees, even though his having the ability and the opportunity to see did not ensure this on their own. But this is precisely what needed to be made intelligible.

What more Hyman does say about these cases (or a subset of them) draws on the idea of "looking or straining to see," which he considers voluntary action. So this analysis of seeing also seems to revert to the models suited to phenomena we hoped to distinguish it from. Hyman doesn't seem to be suggesting that in certain cases, seeing something is a

choice: he says that seeing can *result from* voluntary action, rather than *constitute* a voluntary action. But nothing in this passage offers a way to think about how seeing can *be* a result, other than via the causation accounted for by natural powers.

If these ways of describing what happens when we see turn out to collapse into descriptions of the manifestations of natural powers, then it becomes implausible to say that we have found a way of understanding sight as a distinctive kind of power, or that we have really made sense of it as the potential that is realized in conjunction with a distinctive kind of opportunity. We can say that it is distinctive insofar as its exercises can have either of two different kinds of etiology: sometimes we see things because they cause us to see them; sometimes we see things because something we choose to do causes us to see them. But this account has been spelled out in terms of just two kinds of capacity: the capacity to be affected by things' natural powers, and the capacity to make a choice. It does not invoke a third kind of explanatory structure, and thus does not provide hope that we can make sense of the conception of seeing originally advocated.

To get this kind of opportunity back in view, and to bring it into focus, we need to find a genuinely different kind of explanatory model to place it in. We need to take care in making explicit what is conceptually required, beyond ability and opportunity, for it to be the case that someone sees something. That requires making sense of an opportunity that is not an opportunity to choose. As long as our conception of opportunity only makes room for opportunities to exercise two-way powers, then it will make intelligible only those which present us with choices.[7] And if we are sensitive to the epistemological problems that ensue from taking perception to be a choice, it will seem impossible to make sense of the idea of an opportunity to see, and accepting the causal theory of perception will seem to be obligatory. To say that sight comes with opportunities to see implies that seeing is a matter of *taking* those opportunities, and that will seem to suggest that seeing what we see is something we *choose* to do. If we can't find a way around this, then we will have reason to give up on the idea that there are opportunities to see in the relevant sense, and to accept instead that there are circumstances in which seeing occurs, as a result of a visible object affecting a viewer in a way that causally necessitates it.

I think an alternative can be found – that, indeed, abstention is not the only route to recognizing a discrepancy between getting an opportunity and realizing it. Since *taking* implies the option of forgoing, then we should talk instead about *exploiting* opportunities to see. The alternative to this kind of manifestation of ability is *failing*, rather than refraining. That is, it is the possibility of failure that ultimately allows the concept of opportunity to get a grip. Seeing something is a different kind of success than carrying out one's decision to swim, but it is an "exploit" nonetheless. Of course, simply turning to new words (and even italicizing them) isn't enough to explain what this framework amounts to. In the next section, I'll work on that by pursuing the issue of salience. Hyman is right that this notion provides a way to explicate the sort of phenomenon that happens when we see. But it serves a different purpose than the one he suggests, which I'll explain by invoking its relation to the concept of skill.

4. Salience and skill

Hyman is right to note that salience plays *some* role in occurrences of seeing. It makes sense to reply to a question about why someone saw something that the thing in question was

[7] This issue is related to the debate between Dreyfus and McDowell about the role of the intellect in unreflective action (see Schear 2013).

remarkably salient, or extremely noticeable. But that doesn't mean that the person *could not but see it*. We should not let the relevance of salience convince us that sometimes there is no "distance" between the potential of the overall "can" and its realization. As Nelson Goodman puts it: "No feature of anything is so central or potentially prominent as to not be overlooked even under close or repeated scrutiny" (Goodman 1978, 39). But this negative point leads us to further questions. What role does salience play in the event of someone's seeing, or failing to? What kind of clarification of the situation do we accomplish in referring to it?

Two clues are offered by Hyman's brief remarks about salience. One is the observation that salience comes in degrees (Hyman 1994, 246). If we are going to take this point seriously, then we need to think about the overall structure of the scale on which those degrees lie – what their differences represent, and what the extremes lying at either end are like. His second remark is that assessing the salience of visible things requires taking into consideration facts about "the observer, background conditions of observation, and distance" (Hyman 1994, 246). I will take each of these in turn, and develop them into a different explanation of what is involved whenever someone sees something.

The first observation connects salience, as a way of characterizing a given opportunity to see something, with the notion of skill as a way of characterizing someone's ability. Given that salience comes in degrees, we need to ask: what is the nature of this scale – what kind of differences are differences in degrees of salience? The point that when things are not very salient, a viewer must "strain to see" them provides a clue: if extreme salience shades into a requirement that one to strain to see, then the "salience scale" must be the same as the scale of how hard one must try in order to see a given thing. Degrees of this scale, it seems, represent degrees of *difficulty or ease*: they mark differences in how hard it is to see the thing in question. Something's being extremely salient is a matter of its being very easy to see. Relative non-salience, on the other hand, amounts to being quite difficult to see.

This is related to a point Hyman makes in passing: "the greater one's ability the less one needs in the way of an opportunity" (Hyman 1994, 240). But it reinforces the point that he seems dangerously close to giving up, that seeing differs from the manifestation of natural powers in having a different kind of etiology in *all* cases, not just some of them. Though the effort required to see a given thing might be vanishingly little, and barely worth mentioning, this doesn't mean we should think that seeing could occur with no exertion whatsoever. Moreover, it's not clear why looking at something relatively effortlessly should be understood as its causing one to see it any more than looking hard at it should seem to amount to this. It's here, I think, that the problem with the idea that visibility and sight are "natural powers" becomes most clear, along with the point that the ability and opportunity analysis fits it. The variability in visual salience implies that seeing in general involves effort or exertion that is to some degree skillful. That is, we can't make sense of *absolute* salience, or seeing without looking at all. Recognizing the exertion of looking and the fact that we can look at things in different ways, however, does not entail that we *choose to see what we see*. Rather, we take it upon ourselves to look, and to look in a particular way, and in doing so we see whatever is there to see – whatever we thereby have a view of.

The second observation has to do with how something's salience is to be determined or described. It opens a door understanding the assessment of something's visibility as a rich and varied enterprise, which need not take the form of making statistical predictions about when things are likely to be seen, and can implicate all sorts of features of things, people, and their situations.

When Hyman notes that "salience is relative to the observer, background, conditions of observation, and distance," he gives a few illustrations of this point: "the sound of one's own name is especially salient; the report of a rifle may be salient, but not if it coincides with the explosion of a bomb nearby; white objects are especially salient in ultraviolet light; and what is salient at five yards might not be salient at fifty" (Hyman 1994, 246). This point brings out the way that the analysis of seeing as the realization of opportunity underwrites a certain kind of phenomenological program for describing episodes of seeing – for saying what it is like to have the experience of seeing various things in various circumstances. It allows us to think of circumstances as ones in which seeing a given thing will be challenging or trivial, as a result of the way they are constituted. The way we understand these challenges implicates a complementary understanding of the different ways in which a person can possess the ability to see, as a result of the various kinds of development that this innate capacity can undergo.

When Hyman first makes the suggestion that different sensory modalities have different "opportunity conditions," he says that "in the case of vision, for example, these circumstances are darkness, distance, and occlusion" (Hyman 1994, 241). While these may be basic or fundamental insofar as we must assume that they play a role in accounting for all occasions on which something is seen, that doesn't mean that they are the only features that may be factors, or that any others we might cite can be defined or redescribed in these terms.

In any event, noticing the relationship between these factors and salience makes clear how complex and various their roles can be. Hyman's example of looking at things under an ultraviolet light demonstrates that a thing's visibility doesn't just depend on whether its illumination is sufficiently *bright* – this is just one aspect of the way it is lit, and others will have different kinds of impacts on its salience as well. Moreover, the difference these features make to salience will depend on the nature of the thing to be seen: it isn't that ultraviolet light makes *anything* more salient, but rather that it makes white things more salient than the non-white things they may be surrounded by. The same goes for distance: whether a thing is salient at five, fifty, or any number of yards is partly dependent on its size, for example – or its shape. The triangular shape of a yield sign is more salient at a certain distance than the octagonal shape of a comparably sized stop sign at that distance. Hyman considers at one point whether mist and haze are cases of occlusion, but leaves it with the fact that thinking of them in this way "makes things simpler" (Hyman 1994, 241). But asking the question invites others, whose full answers will be complex. Mist and haze are examples of relative opacity, and it is obvious that relative opacity or transparency bears some intimate conceptual connection with occlusion. But occlusion also bears intimate and interesting connections to relative darkness or illumination (materials can be sheer, insofar as things can be seen through them more easily in brighter light), distance (the water in a murky pond might occlude things lying at the bottom of it, but not those floating closer to the surface), and color (if the pond water has a green tinge, it will be harder to see green things in it than red ones). Camouflage seems to present a visual obstacle that is strikingly different from the canonical examples of opacity and darkness. A purple bead may be easy to see if it is surrounded by green beads, but it may be very difficult to make out if it is surrounded (though not covered up) by other things of a similar purple.

The important thing to note about all of these examples is that they are not cases of being *categorically* prevented from seeing things. These aren't features which *rob* us of opportunities to see things altogether. Rather, they characterize circumstances in which things can be seen, but are relatively *hard* to see. Equally, they characterize what it takes to see something, on the part of a would-be viewer. Thus they illustrate that in general, we

understand the way that different features of a given circumstance contribute to its constituting an opportunity to see by taking them to determine the salience of what there is to see in that circumstance. This means that when we describe opportunities to see, and distinguish them from one another, we are also making it possible to *evaluate* them. We take it that certain parameters make for better and worse opportunities to see certain things, precisely insofar as they determine how easy it is to see these things, or how *well* they can be seen.

So Hyman's point that salience is relative to the viewer in question not only clarifies further the relationship between opportunity and ability, it also opens the door to appreciating viewers' own contribution to the variety of possible episodes of seeing. Hyman only mentions an auditory example of the kind of relevance the perceiver can have: one's own name, when spoken aloud, is more salient than other things one can hear. But it isn't hard to find visual analogues of this. One's friends' faces "pop out" when one scans a crowd; one tends to notice cars of the same make as one's own. However, the observation that salience is a matter of how easy something is to see brings out a different kind of idiosyncrasy of a viewer that can make a difference to whether one sees something. How easy something is to see depends in part on how *good* one is at seeing that kind of thing. This forces a complication of our understanding of the contribution of one's ability to see to the fact of one's seeing something. It isn't just a matter of not being *in*capacitated, but also of what kind of visual capacity one has. It shows that the ability to see isn't a monolithic feature, but one that comes in different degrees and forms: different people, both in possession of the ability to see, may differ when it comes to which types of things they are able to see, as well as in how good they are at seeing a thing of a given type.

Visual salience, as a metric for opportunities to see, coordinates the various features of our circumstances that make it constitute the opportunity – the visible scene – that it is. And it specifies how much exertion will be required of a viewer with a certain set of visual skills, if she is going to see the visible thing in question, in the given context. This makes sense of the fact that, as Hyman says, philosophers of perception have "curiously neglected" the concept of salience (Hyman 1994, 246). While the causal chain conception is in force, there is no need to pay attention to the concept of salience. It is a concept for characterizing a kind of opportunity which is nowhere in sight (pun genuinely unintended).

Once this alternative is sharpened, we can see that salience does not "bridge the gap" between the overall possibility of someone's seeing something, and their actually seeing it. Instead, it could be thought of as a matter of the "shape and size" of a gap, a determination what there *is* to bridge – the kind of skill that a viewer will thereby be exhibiting if she does, in fact, see something that can be seen. What makes the difference, then, is neither the impact of the object, nor the making of a choice, but rather the viewer's success in deploying her visual skill. She must *look* at what's there, and in doing so exercise skill of the right kind and of a sufficient degree. Seeing requires actually looking in the right way, not just being *equipped* to when the opportunity arises.[8]

But we need to be careful in talking about gap-bridging. We should not think of the viewer's effort as an additional ingredient added to the mix of standing ability and encountered opportunity, which explains her seeing insofar as it *makes* her see. Rather, the viewer's exertion of effort simply *amounts to* her seeing, when it succeeds. Nothing can *guarantee* that one *does* look in the way that is required for seeing what can be seen. This is to

[8] For an extremely subtle discussion of the relationship between seeing and different senses of "looking," see Sibley 1955.

reiterate the point of the quote from Goodman, the idea that there is no such thing as absolute salience. It is also the point that like any sense, vision is fallible: visual confusion, or illusion, is always a possibility. And now we are in a position to see, I hope, that this is also the point that Hyman asks us to appreciate when he says that the conjunction of ability and opportunity entail an "all-in 'can'" – the obtaining of a further potential, and not the actual fact of its realization. Its realization is never just an inevitability.

Chapter 2
Seeing How Things Look: Opacity and the Challenges of Vision

If the philosophy of perception is obliged to deal with the business of seeing at the most general level, it might seem as if there is nothing for it to do. A list of things we ordinarily take to be visible can look like an absolute hodgepodge – even a mundane a list of things a person took herself to have seen on a given day might include a collection as diverse as the living room, the armchair, the neighbor's cat, her reflection in the mirror, Mount Tamalpais, the sidewalk, the sunset. Can we expect to identify any particular unity that these things exhibit, which makes them candidates for such a list? Or is this expectation unreasonable and doomed to leave us dissatisfied?

This chapter suggests that, despite this heterogeneity of the visible world, there is philosophical work to be done in the interest of a clearer understanding of our capacity to see. It approaches this by noting that despite the variety of its manifestations, there is a feature that seems to characterize any paradigmatic occasion of seeing, and that is opacity. When people see things, they are in the vicinity of something opaque.[9] Most of the examples on the list above refer to opaque things explicitly, and the ones that don't – the living room, the reflection in the mirror, and the sunset – refer to situations that would contain something that is opaque. Considering the challenge that opaque things present to us as viewers will point us to the importance of a notion of visual skill, or the way that our power to see can be developed. It will also lead to clarification of the role that things' *looks* play in the business of meeting this challenge and related ones.

I will work toward this conception of seeing in a somewhat oblique way, by examining some worries that lead to the conclusion that we must reject what we are ordinarily ready to say about what we see, and accept in its place a humbled conception of our capacity for sight. These worries are rooted in the facts of opacity, but in these discussions, opacity usually remains an unspoken factor while its conceptual partner, the notion of occlusion, is invoked to construe the challenge of seeing opaque things as one that would take a miracle to surmount. By reflecting carefully on the hallmarks of a visible scene and the circumstances of confronting opaque things, we can assuage these worries and replace them with a better way of thinking about the phenomenon of opacity.

I start by considering lines of thought that Thompson Clarke presents in "Seeing Surfaces and Physical Objects," which generate puzzlement about our ordinary conviction that we are able to see things like tables and chairs by making tacit use of the fact that they are opaque objects. I argue that these lines of thought are faulty, and that the humility that they evince about what our sight allows us to see result in is undue. I then consider work by Sean Kelly and Alva Noë, both of whom delve into our ways of dealing with our circumstances in order to see things in our surroundings. I argue that their positions exhibit a different kind of humility about our capacity to see – one that is also unwarranted, and is related to an undue confidence about our ability to see how things *look*.

1. An undue humility about what we can see

Thompson Clarke's paper "Seeing Surfaces and Physical Objects" provides a revealing illustration of how the phenomenon of opacity gets problematized. Clarke presents the reasoning behind a revisionary way of thinking about seeing by focusing on a hypothetical

[9] This is another way of putting Noë's point that we can't make sense of the idea of total x-ray vision (Noë 2012, 42).

encounter with a tomato. First, he establishes the scene by asking us to "[c]onsider a situation in which another person X is looking at a physical object, say a tomato, in normal conditions… Perhaps he is seated at his desk in his study and the tomato is in front of him on his desk." Clarke acknowledges that the tomato he has described fits the paradigm of a visible object: "this is the sort of situation in which we'd say that another person, X, can see a tomato." From there, he immediately goes on to note some further implications of the description of the scenario. "But," he says, indicating that these implications are meant to put pressure on the conviction that the tomato in this scene is visible, "how much of the tomato can X see? He can't see any of the far side or inside…All he can actually see of the physical object is a portion of the surface towards him…" Clarke then makes the conflict explicit: "We said that this was the sort of situation in which we'd say that X could see the tomato. But can he? No, he can't. Obviously all he can *really* see is that portion of the surface towards him…So now we see that we were wrong to suppose that X could really see the physical object itself" (Clarke 1965, 99-100).

In this line of reasoning, asking "how much of the tomato" a person in these hypothetical circumstances is able to see turns out not to be so innocent a question: it forces us to admit that he can see only "that portion of the surface towards him," and cannot see the tomato whose surface it is. The question is answered by taking note of certain facts about what the person in the scenario can't see: he can't see "any of the far side or inside" of the tomato. And this is taken to show that what he can see is nothing more than and nothing other than a *part* or *aspect* of the tomato: a particular region of its surface.

But does it show this? Clarke's own assessment is that the investigation doesn't *show* that this is the case, it rather *makes* it the case, by shifting the terms in which an answer to the question "what can someone in X's position see?" must be given – from terms like "the tomato" to terms like "this part of the tomato's surface." He uses this assessment to reject two distinct philosophical positions: one, that all we ever see in looking at things like tomatoes[10] is parts of their surfaces; the other, that what it is to see a thing like a tomato just is to see a part of its surface. The first position is mistaken because we have no reason to deny that there are contexts where "a tomato" is an appropriate term for answering the question "what can someone in X's position see?" The second position is mistaken because in seeing a part of a tomato's surface one's "perceptual position" is "worse off." It is the same position one would be in if one were faced with a mere *chunk* of a tomato, which included that portion of the tomato's surface (Clarke 1965, 101-103).

The rejection of both of these positions may ultimately be correct. They are both opposed by strong convictions: one, that we sometimes see things like tomatoes, and two, that facing a tomato puts one a different "perceptual position" than facing a tomato chunk. But Clarke's reasons for rejecting them are no less puzzling than the motivation for accepting them. Clarke's analysis asks us to accept something that seems just as contrary to common sense as his opponents' conclusions: that the question "how much of the tomato can be seen by someone in X's position?" is somehow nefarious, that it has the power to downgrade a person's perceptual position, rather than prompting, as an "innocent" question would, a genuinely fruitful inquiry. Why can't we answer this question simply, straightforwardly, in a way that doesn't inherently conflict with any of the convictions we start out with when the question arises? I think we can.

[10] Clarke identifies the relevant class as that of "physical objects," but as will become clear, the point in fact generalizes to opaque things specifically.

Let's look again at the question and the way the investigation proceeds: "how much of the tomato can X see? He can't see any of the far side or inside…All he can actually see of the physical object is a portion of the surface towards him…" Focus on the claim that the viewer can't see the far side or inside of the tomato. What is it that supports that assertion? How is it that in understanding the nature of scenario as described so far, we understand this to be true of it? I suggest that it is because we take the inside and the backside of the tomato it mentions to be blocked, or occluded, from X's view. And I suggest that we do so because we take the tomato itself to be an opaque object, an object which you can't see through. That consideration is what underwrites the claim that the viewer in the scenario can't see the particular parts or aspects of the tomato.

But what should we say is blocking X's view of these tomato parts, exactly? This turns out to be a bit of a delicate matter – though not impossible to address coherently. It is natural to say that the *front* part of the tomato occludes the rest of it. But what exactly is the front part of the tomato? Is it some portion of *volume* of the tomato? The troubling thing about this answer is that if we take this volume to be opaque, then it too will have an inside and a backside that are occluded from view, and so there must be some smaller volume that is sufficient for the occlusion of the hidden parts of the tomato. We might then conclude that it is not a portion of the tomato's volume, but rather a portion of its surface, that accomplishes the occlusion of its other parts. Talk of "opaque surfaces," which is relatively common in philosophical discussions of vision, expresses at least a vague or tacit commitment to this way of thinking. But it is doubtful that it is correct, when we realize that only things that have some volume can be called opaque, and thus only things with some volume can accomplish the relevant kind of occlusion.

In fact, it is misleading to think of the tomato's front as either a portion of its volume or a portion of its surface. The tomato's having a front is not a matter of its comprising any particular chunk of vegetable matter. Rather, it is a matter of the tomato's being *oriented* in a particular way, in this case, relative to a viewer's gaze. Distinguishing the back of the tomato from its front does not tell you anything about the kinds of parts that make it up (as distinguishing the stem, seeds, pulp, and skin would). Instead, it tells you what direction things lie in from the tomato, or how the tomato is situated relative to other things. That *this* is the front means, in this case, that the viewer is facing the tomato from *there*. So although the front of the tomato can be indicated by delineating or making salient either a portion of its volume or its surface, that does not seem to be what is ultimately indicated by the demonstration.

Given all this, we might instead say that it is the tomato itself, rather than any particular part of it, that occludes its inside and backside. Opaque objects, we can say, are essentially self-occluding. In ending up here, however, we must be careful not to lose sight of two closely related things. First, the self-occlusion that an opaque object exhibits is occlusion of some of its own *parts*, and *not* occlusion of its entirety, of itself as a whole. Second, we must acknowledge a status that the facing surface, or front, of the tomato has which the other parts do not have: it is *not* occluded, and thus it *can* be seen by X, while the other parts cannot be.

If all this is coherent, it seems we can, after all, regard the question "how much of the tomato can X see?" as a philosophically "innocent" one. The reflection it prompts can be carried out on the basis of common sense – though Clarke's term "surface inquiry" is something of a misnomer, and his reconstruction of it is dangerously quick. The result ("the front, but not the back or the inside"), moreover, in no way threatens our ordinary conviction that someone in X's position can see the tomato. We can hold on to that, while

admitting both that the facing surface is visible and its back and innards are not, and that seeing its facing surface and seeing the tomato itself are not equivalent "perceptual positions." We can maintain that they are different things to see, though both *can* be seen by the viewer in this scenario, because, relative to his gaze, neither is *on* the far side of an opaque object, and neither *is* the far side of an opaque object.

So much for Clarke's suspicion of the "how much" question and the "surface inquiry" that it prompts. The latter can be carried out common sensically, and so the former can claim its philosophical innocence. However, there is more to say. The phenomenon of opacity can still be a source of puzzlement about the role that various types of things can play in our visual lives. To see how, we must shift from considering opacity's role in occlusion, to considering its consequences for variation and similarity in the ways things look.

2. An undue humility about how well we can see things

A different tradition of philosophical thinking on perception gives a take on the "how much?" question Clarke poses that provides an interesting and instructive contrast. In certain ways, it occupies a middle ground between Clarke's two opponents. It rejects the position that our vision stops at the surfaces of opaque things, barring us from ever seeing these things as wholes. But it also rejects the nonchalance of the position that simply equates seeing a thing with seeing a part of its surface. Instead, it holds, seeing an object is a matter of seeing its surface *in a particular way*. The relevant way takes some care to spell out, but it involves taking into account or being sensitive to ("observing," somewhat in the sense in which one observes a holiday) the way the surface contributes to the thing as a whole.

As representative of this line of thought, I'll consider work by Sean Kelly and Alva Noë. As inheritors of ideas from phenomenologists like Maurice Merleau-Ponty about the role of the body in perception, both reflect carefully on the way in which we are situated in the world we perceive, and both emphasize that perception involves skillful activity on the part of the subject. By my lights, they have done as much as any thinker to show that these ideas promise to elucidate philosophical issues about perception and its role in our conscious lives. Because they find different ways of articulating the way we should put them to use, it is rewarding to contemplate and compare them. Both of them, I think, give accounts whose acceptance requires an undue humility about our ability to see opaque objects. In this section, I bring this out in Kelly's discussion in "Seeing Things in Merleau-Ponty" and in Noë's enactive view as elaborated in *Action in Perception*. In the section that follows, I uncover the flipside to this kind of undue humility about our capacity to see opaque things": an undue confidence in our ability to see how things look.

Though Kelly's paper is a work of interpretation, it consists in large part of independent reasoning about selected insights found in Merleau-Ponty, because Kelly thinks that in this area Merleau-Ponty "didn't quite get his own view right" (Kelly 2005, 76). Here I will be concerned with assessing the views that Kelly thinks Merleau-Ponty had in mind on their own merits, leaving aside the question of whether they provide a faithful interpretation.[11]

On the way of thinking that Merleau-Ponty's work ultimately inspires, says Kelly, we have a general capacity to see opaque things, but our view of them is always relatively poor. A "full view" of such an object is not one we can ever have. The world is visible and we do manage to see, but the opacity of certain things in it is ultimately an obstacle to our seeing. It

[11] For a discussion of whether Kelly successfully interprets Merleau-Ponty, see Matherne 2017.

is an obstacle that we are in general able to overcome, but we can overcome it only partially. Thus every view we get of an opaque thing is suboptimal, every experience of seeing a given thing deviates from the ideal. Whenever we see an opaque thing, some of it always remains "hidden" from us, and in this way what we see "transcends" or "goes beyond" our experience of seeing it (Kelly 2005, 76). This is true for opaque objects because of the self-occlusion opacity entails, but the basic phenomenon is more general. The same thing goes for things like tables and chairs and mountains, as well as features like their colors, shapes, and sizes: any experience of seeing such things is sub-optimal in some way. All objects of vision have their own way of "hiding" themselves, their own self-concealing nature.

Interestingly, Kelly thinks of the case of seeing a thing's color as illustrating this point particularly clearly. He considers it a "test case" for the account, on which one can model an appropriate treatment of phenomenon of seeing opaque things (Kelly 2005, 79). But Kelly's brand of humility, his claim that we only ever get sub-par views of opaque things, is an immediate consequence of his conception of the nature of an ideal view. Kelly has it that an ideal experience of seeing an opaque object is one in which the object is "seen from everywhere all at once" (Kelly 2005, 91). It is meant to be clear that this is not a view we can ever have. Whenever we perceive (whenever we do anything), we are situated somewhere in particular, and we occupy a certain spatial relation to things in some environment. Other situations, other perspectives, are accessible to us in principle – our current position is only one of many that we could come to occupy. But from where we start, coming to occupy them is possible only through movement, and movement takes time. So we occupy different vantage points only in succession, never all at once.

Kelly claims that this conception of the optimal viewing conditions for an opaque object is "motivated by a genuine insight…that no single point of view reveals the object fully" (Kelly 2005, 91). But this is too close to the claim about what the optimal conditions are for it to provide much motivation for it. If the position has any plausibility, it is to be found in the observation that there are better and worse views of a thing – better and worse circumstances relative to seeing it, visually picking it out from its surroundings (as well as seeing what it's like, visually determining the features it has). Kelly introduces this idea in advocating the view that there are "inherently normative, rather than descriptive, features of visual experience" (Kelly 2005, 87). Though some have objected to the idea that seeing is normative, the observation I just mentioned should not be regarded as a controversial or surprising philosophical position. On the contrary, it is implied in our everyday talk about the things we see. We easily make, understand, and evaluate claims about whether, on some occasion, someone saw something well, or poorly, whether their view of a thing was a good one, or whether it was compromised in some way. We readily apply norms or standards to cases of seeing in this way.

However, our everyday talk does not speak to the conception of the relevant norms that Kelly ends up suggesting. If asked to spell out the optimal conditions for seeing a particular opaque object, we would have to make use of the following rubric: close enough but not too close, surrounded by a sufficiently transparent medium, without any opaque objects in the way, amongst sufficiently contrasting objects, without and in sufficiently bright but not blinding lighting. What these constraints turn out to be will depend on the specificities of the thing in question: its color, shape, size, and so on. (A good view of the table's color, shape, or size may correspond to a different set of parameters – more on that later.) For example, from a distance that deviates from the ideal parameters significantly, you will get a poor view of the thing in question: you may see it, but not very well. Too much

more distance and the thing will be out of view. In general, circumstances that afford bad views shade off into ones that don't count as views at all.

Kelly's conception of the ideal as the "view from everywhere all at once" doesn't fit with these commitments. It contradicts the way that we take views to count as deviations from the ideal. According to him, no matter where I stand when looking at an object, "the view from everywhere is an ideal from which I can sense myself to be deviating" (Kelly 2005, 91). But while it is easy to suppose that in a particular circumstance I am too far from an object to see it well, and I sense that to see it better I'd need to be closer, I wouldn't make sense if I claimed that I can't see it well because I'm too *here*, or too *somewhere*, not *everywhere* enough. This is not the way that bad views compare to good ones. Just as bad views shade off into the absence of any view, they also shade into good ones – not just "better" ones, but ones that we consider ideal, in which one is able to see the thing in question perfectly well. For these reasons, it seems to me that Kelly's construal of the circumstances in which we can see something well is untenable. The claim that we are never in a position to see an opaque object fully expresses an unwarranted humility about our capacity to see. Given that, we can't accept the general point about perception that he expresses in saying that we "experience objects as *transcending*, or going beyond, [our] experience of them" Kelly 2005, 76).

In the interest of identifying truths in the vicinity (and distinguishing them from tempting mistakes), it is worth comparing Kelly's view to the "enactivist" account presented by Noë in *Action in Perception*. Though Noë and Kelly dispute certain points, I am here concerned with their affinities. The way in which Noë accepts the kind of humility I've been considering comes closest to the surface in his claim that perception is "virtual, *all the way in*" (Noë 2004, 134). What he means by this is that visible things essentially "outstrip" our ability to be visually aware of them.

Noë bases this claim on considerations about the visibility of things' aspects – the same considerations that motivate Kelly and Clarke. Any thing, he notes, as well as any of its qualities, can appear differently in different circumstances, and so any particular encounter with a it will be an exposure to only one of these many possibilities and no more. In this quote he applies the point to color, shape, and the things that have them:

> "Just as it is not possible to see every aspect of an object from a single vantage point, so it is not possible to experience every aspect of an object's color all at once, from a single vantage point (as it were). The color of the object is no more completely visible under a single set of viewing conditions than is the shape of the object visible from a single vantage point. (Nor is it less visible.) How the thing looks with respect to color in these conditions is not enough to tell its color." (Noë 2004, 128)

Noë does not explicitly affirm the principle Clarke considers, that seeing a thing requires seeing all of its parts or aspects; in fact, he explicitly denies that it requires this. But nonetheless, he implies that the *best* view of something, a *completely clear* or *full* view of it, would be one which exposes all of its parts, aspects, features, facets. He does not make this assumption explicit like Kelly does, in the form of a claim about the ideal or optimal conditions for viewing any particular things. But the assumption provides the transition between the first two sentences in the quote above. The first sentence makes a claim about whether a thing's aspects are visible from a particular vantage point – not all of them can be seen. The second makes a claim about how visible the thing *itself* is, from that vantage point – it is not "completely" visible.

So while Noë rejects the first position Clarke reconstructs, and denies that if one can't see all of a thing this entails that one does not see it *at all*, he nonetheless holds on to the conviction that this puts one in a "poorer perceptual situation" than one might hope for. It's not that one sees something else, some less significant item. But one's position is still limited or diminished, in comparison to some alternative. We can see things, when some of them is occluded or otherwise out of view, but we can't see them as well as we would if this were not the case. A circumstance in which we would ordinarily say that someone sees a tomato may be just that, but this view is compromised by the tomato's opacity.

The consequence is that vision only ever enables us to get glimpses of things, to get indications of what they may be (tomatoes) and what they may be like (red, round, and bulgy), but which never count as having them or their features fully, squarely in view. As Noë puts it, "qualities are available in experience as possibilities, as potentialities, but not as givens," and "experience" – the experience of seeing – is "a dynamic process of navigating these possibilities" (Noë 2004, 135). Thus the visibility of a thing is a potential that viewers (or at least, viewers like you) can never fully realize. Every look at something is nothing more than a glance.

Noë's response to a review by Michael Martin, which concentrates on Noë's stance on the principle that seeing a thing requires seeing all of its parts or aspects, sheds some light on what prompts him to take this stance. There Noë rephrases the claim that perception is virtual through and through by saying that "*all* perceptual presence is the presence of access." That is, what characterizes the presence things can have when they are "open to view" is "only the ease of access" – only that "we need to do *less* to modulate our sensory relation to that which is open to view than we need to do to modulate our sensory relation to what is, as it were, hidden" (Noë 2008, 697).

But this formulation belies the relation between what it is for something to be *visible* and what it is for it to be *seen*. Though the latter implies the former, the reverse does not hold: something may be visible to someone, yet not seen by them. The difference between these two circumstances must be kept in sight when we consider questions of *ease*. If something is *easier* for us to see, then we are more likely to see it, but its being easy to see is not the same thing as its actually being seen. Moreover, no matter how easy something is to see, that does not entail that anyone manages to see it.[12] To say that something is there to see is always merely to characterize a possibility or potential. But Noë seems to want to say that the phenomenon of one's actually seeing something is itself a matter of potential – one that takes relatively little to realize.[13] It may be right to characterize a thing's being easy to see as its being relatively "accessible," and thus as having a "virtual presence." And it may be right to pick out this potential by reference to ways we can "modulate our sensory relation" to things. But that does not mean that when the thing *is* seen, this potential is not fully realized.

If the distinction between visibility and being seen is maintained, then to say that perception is virtual through and through, or that all visual presence is a matter of access, can only be to say that while things are visible, they are never actually seen. Noë doesn't intend to say that, however. Instead, he means to say that the visibility of a given thing is a potential that viewers only ever realize partially. So, like Kelly, Noë is moved by reflections on opacity and analogous phenomena to offer a "humbled" view of visual perception, even though he rejects the first position that Clarke reconstructs.

---

[12] This point is discussed in the previous chapter; see also Hyman 1994.

[13] It is worth noting that the conflation involves the ideas of two different visible things. In speaking of the ease of seeing something, Noë has in mind an aspect of a given thing, like the facing surface of a tomato.

In both cases, I have argued, arriving at the view involves relying on a faulty premise or implicit assumption, without which we have no reason to accept any kind of denigration of our ability to see opaque things or their properties. But the fact that they hold on to humility in some form while aiming explicitly to reject the extreme position that Clarke reconstructs calls for further investigation. The pervasiveness of this tendency makes one suspect that there are truths in the vicinity, and temptations to misconstrue them, that have yet to be brought to light. I will work towards this in the rest of this paper, by delving into further aspects of Kelly's and Noë's views. In the next section, I will show how their humility about our ability to see things is bound up with an undue confidence about our ability to see how things look. In the final section, I will distinguish their conception of the way that our capacity for sight is outstripped or transcended from different way of considering our finitude as viewers.

3. Undue confidence, or a reversal of the relative ease of seeing things versus seeing how they look

As noted above, Kelly and Noë both reject the principle that characterizes the position of the first philosopher Clarke imagines, that if you see a thing, you must see "every bit" of it. But they also reject the opponent Clarke imagines, who claims that seeing a thing *just is* seeing *some* of it. For them, just seeing some of a thing isn't sufficient for seeing it. If it is true of a person that she sees something, more must be the case than just that she sees some of the thing. But the further condition is not that she sees the rest as well. Rather, it is that she sees some of it with the understanding, or the knowledge, of the relationship between the aspect she sees and the whole it belongs to. This relationship is that of an object's *look* (in a certain setting) to the thing that has that look (in this setting). For both Noë and Kelly, that is what seeing the thing amounts to. They diverge on certain questions about what this knowledge amounts to, and how we get it. But despite those differences, they conceive of the structure of the event of seeing something in the same way.

Kelly is led to his view by reflecting on the idea that the "most basic unit of perceptual experience is the presentation of a figure against a ground" (Kelly 2005, 96-97). What this means is that whenever we see something, that experience involves something in our surroundings serving as the object of our focus, and other elements of our surroundings remaining in the background, constituting the situation in which we and the focal object are to be found together. Only what is in focus is *seen*, or apparent (Kelly uses the term "determinate" to refer to this status); elements in the background are unseen, being "out of focus" or "hidden" (in Kelly's terms, "indeterminate"). But their relation to us differs from that of other objects we don't see on this occasion, things which are "out of view" in a stronger sense. To count as being in the background, a thing must be part of one's current environment, in contrast with all the other business that the universe contains. The background and the focus together contribute to, or constitute, the view we have of whatever it is we see.[14]

Kelly is concerned to do justice to the way in which things in our surroundings can contribute to our view of something by remaining in the background, unseen. He takes cues

[14] To illustrate: The Eiffel Tower, located 5,550 miles away from me on the Champ de Mars, is analogous to the stack of books underneath the laptop I'm typing on, insofar as both are currently unseen by me. But the books have a hand in shaping my experience of seeing my laptop, with which the Eiffel Tower cannot compete. The placement of the books determines the angle of my gaze and the distance at which I must focus, in having this experience of seeing my laptop.

from certain remarks by Merleau-Ponty, such as the claim that "the perceived contains gaps that are not mere failures to perceive" (quoted in Kelly 2005, 80), and that "lighting and reflection" as well as "shade" serve to "lead our gaze instead of arresting it" and "becom[e] our environment in which we establish ourselves" (quoted in Kelly 2005, 84). Illumination thus provides one kind of example of a background feature, but the category extends to the space on the far side of the opaque objects in front of a viewer (and whatever it contains), as well as the occluded parts of the opaque object they are facing. The task is then to understand how these features can all play analogous roles in characterizing an episode of seeing.

The "rubric" I gave earlier on for the parameters of a good view of a thing indicates a way of understanding this. The features now under consideration are precisely the ones that matter to a determination of the *quality* of the view in question. These aspects of one's situation in relation to an object – the way light shines on it and the way it reflects that light (a matter of its opacity and its color), the distance between you and the object and what intervenes in it, what lies next to and behind it (the proximity of other objects with similar characteristics) – all of this determines how good a view of the thing your present circumstances afford. This tells us something about the kind of "positive role" that the background makes to an experience of seeing something, and points to a way in which an unseen visible thing can represent not a "mere failure to perceive" but a factor in the perception of something else.

But this is not quite what Kelly seems to want to say. Inspired by Merleau-Ponty's talk of light "leading the gaze," Kelly wants to say that the contribution that things in the background make is that of a "positive presence," though their presence is different in kind from that of what is in focus. On his account, an element of the background "present[s] itself not as a determinate quantity but rather in terms of how well it enables me to see the thing I'm looking at" (Kelly 2005, 84). As present in this way, they "tell me something about what should happen for me to get a better, fuller, more complete experience of the focal object" (Kelly 2005, 90).

Kelly illustrates the idea with an example of seeing the color of a tabletop when it is unevenly illuminated, as host to an array of cast shadows. In seeing the tabletop's color, "I see, in a direct bodily manner, *how the light would have to change* for me to see the color *better*" (Kelly 2005, 85). This is not the kind of experience one would have if the lighting were the object of focus, if it were the thing one is currently in a position to *inspect*. It is not an experience of *seeing* the light. But Kelly still considers it a kind of experience of the light. Having this experience makes one informed of the lighting's consequences for the quality of one's view. So on this account, the aspects of one's surroundings that provide the background in one's view play a normative role not only insofar as they dictate how good one's view is, but also insofar as they are *prescriptive*: they inform the perceiver of the quality of her opportunity to see the thing in question, and they dictate *to her* what must happen for her to exploit it. Kelly seems to infer this from a need to account for the fact that if you decide to look and see what color the tabletop is (if you are "asked to determine the color of the table"), you make the adjustments necessary to get the color in view: "your eyes move automatically to the part of the surface where the lighting is best" (Kelly 2005, 84).

Kelly says this kind of experience of illumination is "at the same time" an experience of the tabletop's color – an experience of "seeing the color to be one shade rather than another" (Kelly 2005, 86). He draws a conclusion from this that he admits is surprising, that "the real color is *never* determinately seen" (Kelly 2005, 87). He explains: "The reason for this is that even when the lighting conditions are optimal, they are still experienced as a deviation

from a norm, only in this case the current lighting is experienced as a 'null' deviation from the norm" (Kelly 2005, 87). This is indeed a "surprising result," since "even if the lighting is not experienced as a determinate quantity, you might have thought that the color it illuminates could nevertheless be experienced as a determinate shade. Because of the way figure and ground are interrelated, however, this simple view cannot make sense" (Kelly 2005, 86).

In drawing the conclusion he does about this case, Kelly prepares the ground for the claim that the optimal view for seeing an opaque thing is the view from everywhere, a view we can never actually have. It has us think that the possibility of assessing occasions on which a given thing is seen entails that any view of a thing we can ever hope to get will fail to allow us to see it fully clearly. If this is the case, then the fact that a conception of the optimal conditions for seeing something makes a good view unattainable is not a mark against it.

However, the point that the experience of the illumination is at the same time an experience of the object's color does not in fact have this surprising consequence. There is no barrier to counting the experience as both a determination of the color of the tabletop – an experience of seeing its real color – *and* as an experience of being guided by the quality of the light in doing so. Determinacy and indeterminacy can both be exhibited by a single experience – indeed, this possibility is presupposed by the point that every experience of seeing something is characterized by an object of focus and a background against which it is distinguished.

But Kelly does not seem to allow this. So, having noted that illumination is a factor in the quality of a view of a thing's color, Kelly goes on to affirm that if one *gets* the view that these circumstances afford, this must involve *seeing* something other than the color, namely the quality of the illumination. This is where looks get into the picture, so to speak: Kelly identifies the quality of the illumination with the way the color looks (in these circumstances).

He does so because in discussing the example of seeing the color of a tabletop when shadows are cast across it, he observes that the illumination of the tabletop in question could vary, without necessarily making it the case that the tabletop's color can no longer be seen. Its color, that is, is visible in a range of circumstances that differ with respect to the illumination. Philosophers often discuss this general phenomenon under the heading "perceptual constancy," since it involves a thing or feature remaining constant as the object of focus in a range of visual experiences whose other parameters vary. Philosophers tend to be as concerned with the way these various experiences differ as they are with the fact that something remains constant across them. Kelly points out what is usually common ground, that these variations affect the way things *look* (in this case, the way the tabletop's color looks) (Kelly 2005, 84).

According to Kelly, in seeing the way things look, we are informed of our circumstances, and thus we are "ready to deal" with what is hidden in the background (Kelly 2005, 100) and can "automatically" make the adjustments necessary to achieve and maintain the best view of the thing that we can get.

When it comes to how we manage to see how things look in the first place, the answer seems to be quite a different story than the one for things like tomatoes and their parts and features. Seeing how things look seems to be a comparatively trivial enterprise: the environment simply *shows* how it looks to us, thereby showing us what it already "knows and sees" – and what we must do to see what's in it (Kelly 2005, 102). We ordinarily have a

perfectly clear view of how things look, which can give us only a compromised view of the things that have these looks.

Though it seems right to say that seeing things requires making adjustments and dealing with our situation in the ways Kelly points out, I do not see why we should accept this account of how we do this. I would not deny that the experience of seeing a thing counts as an experience *of* the circumstances of its situation (not merely as an experience *in* those circumstances, nor merely as one that is had *while* those circumstances obtain). Nor would I deny that it is an experience of their quality for seeing the thing in question. But the point that the factors that determine how good a view is can also account for how we are able to see a given thing does not entail that we only ever see their quality clearly, and never get the chance to see them perfectly well. The fact that we are able to see something because of the quality of our circumstances does not entail that we *see* that quality *at all*. An experience of seeing a given thing can be counted as an experience of the background against which it is distinguished because it is an experience of *coping* with them successfully, in the way required given their quality – coping with them in the way required for seeing *something else*. And this coping need not involve seeing how things look.

Acknowledging the fact that we automatically, readily, skillfully do things that enable us to see things that are there to see need not lead us to accept that we are prepared for this by our present circumstances, in seeing how the things before us look. That readiness can be chalked up to the general ability that we bring to the challenges that our present circumstances constitute. This ability is something we have acquired from *past* experience – encounters with *other* circumstances have prepared us to face new ones. And when we see what there is to see now, this is not predicated on any prior awareness of the peculiarities of our current situation, but rather on a standing ability to achieve a focus on things like the one we now confront, in circumstances like the present ones.

But the heart of the problem is not in Kelly's specific conception of the looks of things, as elements of our surroundings that show us, on the fly, how to look the things that have them. This can be seen in the fact that Noë's "enactive account" that Noë contains a different conception of the looks of things and their role in episodes of seeing, but it too leads to the thought that we can't see opaque objects (or their shapes and colors) very well, relatively speaking. I turn to Noë's account now, to make clear the more general issue of the relationship between getting a view of a particular thing, and getting a view of the way that thing looks.

As noted above, Noë does not invoke the notion of an unobtainable but ideal view from everywhere all at once. Yet he is led to the commitment that perception is virtual through and through by a view that aims to account for the same observations that animate Clarke's and Kelly's discussion. The approach he settles on invokes the notion of a different kind of "background" than Kelly's: instead of supposing that our surroundings inform us of how to get things better in view, Noë considers what we as viewers bring to the table, so to speak, no matter what circumstances we end up in. It is because we have certain skills, a certain kind of understanding of the world's visibility, that we manage to see the things around us. So Noë diverges from Kelly in attributing knowledge of what it takes to see only to viewers, refraining from attributing it (and the ability to impart it to us) to the circumstances in which things can be seen. But beyond that difference, their respective conceptions of what is ultimately involved in seeing something display significant affinity.

Noë's account has it that your perception of an opaque thing before you "depends on your tacit understanding of the ways its appearance (how it looks) depends on movement" (Noë 2004, 77). Noë calls this understanding "sensorimotor knowledge" or

"sensorimotor skill." On his analysis, seeing a thing is a matter of seeing how it looks – experiencing some visual stimulation that it affords – and taking it to be part of a larger set of possibilities, each of which correlates with some potential action one could take to adjust one's situation. The experience of seeing a tomato, then, is the experience of seeing how it "merely" looks, plus understanding the pattern of that this particular look fits into. The same goes for the features of a tomato – its size, shape, color, and so on. For anything visible, we see it by seeing how it looks. Seeing how a particular thing looks serves as the basis of our success in seeing *it*, if we do.

Thus, Noë and Kelly both are motivated by a conviction that in general, seeing something implies seeing how it looks, because seeing how a thing looks serves as the basis for seeing the thing itself. A form of this conviction seems to be shared by the philosopher Clarke imagines, who raises no issue about the visibility of a thing's surface. So the undue humility about our ability to see opaque things is bound up with a kind of confidence about our ability to see how opaque things look.

That confidence, I now want to argue, is unwarranted. It is not warranted as part of an account of how our views of visible things are sub-par, because the humility that such an account entails is unreasonable. That is, it is not something we need to accept in order to make sense of how we manage to see things themselves, in the way that we do. Moreover, it does not do justice to the distinctive *difficulty* we face if we are interested in seeing how a thing looks. In the remainder of this section, I sketch a way of thinking about that challenge and the skills involved in meeting it.

Even if we reject the idea that things in our environment lead our gaze, or that we see things by seeing how they look, we can still affirm that seeing is something that implicates the possession of skill on the part of the viewing subject, and that the way I cope with my circumstances in order to exploit opportunities to see things is a matter of the conscious exercise of that skill, not mechanical response. And that means that the way this exercise unfolds is available to me for reflection and further investigation. This kind of investigation is precisely what we do, I suggest, when we manage to see how things look.

Seeing how a thing looks isn't something we are necessarily ready to do, or perfectly equipped to do, whenever we look at it. But given that looking at it is something we do consciously, it is always in principle possible for us to extend our visual capacity to include seeing it. That means that the way this exercise unfolds is available to the viewer for reflection and further investigation. This kind of investigation is precisely what we do, I suggest, when we try to see how things look. Seeing how something looks involves reflecting on the nature of the circumstances in which it can be seen, when this reflection takes a visual form. This isn't something we need to do in order to see the things in question. But it is true that it is something we *can* do when we see the things in question. That is, while it isn't the case that when one sees a thing, one necessarily sees how it looks, it is true that one thereby has an *opportunity* to see how it looks. Given that you see the thing, it will look some way to you, and that means that these circumstances will be there for investigating. Investigating them is investigating further, and will involve drawing on further skill. It may require acquiring new skill, developing the skill one already has. Seeing how a thing looks isn't something we are necessarily ready to do, or perfectly equipped to do, whenever we look at it. But given that looking at it is something we do consciously, it is always in principle possible for us to extend our visual capacity in a way that enables us to see it.

If a thing is opaque, seeing its shape may amount to a fairly complex experience. It might require getting a look at a variety of its parts, and some of these parts may be occluded by the object when it is seen from an angle that puts others in view. So getting a view that

enables one to see the thing's shape might involve some maneuvering – perhaps backing up to an appropriate distance and walking around it a bit, perhaps bringing it in close and manipulating it in one's hand. This is a more "expansive view" than what Kelly seems to have in mind when he talks about seeing something from a "single point of view." On the other hand, it is not a view "from everywhere all at once" – it is a view that one gets over the course of some "exploration," which takes some time, and certainly doesn't take one *everywhere*.[15]

None of this entails that the thing that *has* the shape can't be seen perfectly well from a more narrow point of view, the kind Kelly has in mind when he talks about a "single vantage point." The fact that a thing's shape has consequences for what it takes to see it does not entail that one needs a *view of its shape* in order to see *it* (or, to see it as an object of *some* 3D shape). When an opaque object is seen, its shape (along with its color, the way it is illuminated, its distance, its beauty) can remain in the background, accounting for how we are able to see it, but remaining unseen.

4. Vision and conclusion

To conclude, I will summarize some of the lessons that have come out of the discussion above.

Things' colors, shapes, and sizes – their natures, in general – determine the parameters of a good view of them, and constitute aspects of how they look. Seeing an opaque object – and seeing it as three-dimensional or voluminous – should not be conflated with seeing how it is shaped. None of its visible features need be seen, given that the thing itself is seen. A thing's opacity is not so much an obstacle for seeing it, but an enabling condition of seeing it. It provides the structure of a scene in which it is visible, marking a distinction between the thing as an object of focus, and the surroundings that form its background.

If it is true that the things we see "transcend" or "go beyond" our experiences of seeing them, this is not because we can't ever get a full view of them, or a view of them as wholes. It is worth noting that this latter possibility is compatible with two other points. First, that what we see is not identical with our experience of seeing it. What we see may "go beyond" our experience in the sense that it is not "mind-dependent." Second, that our senses are fallible, which entails that what people seem to see comes apart from what there really is (to see).

We should also distinguish between the point *seeing everything there is to see all at once*, and not *seeing a thing fully or perfectly well*. The former is not a possibility, the latter is. When we see, we don't see everything there is to see all at once. We see one thing that is there to see, instead of all the rest. Nonetheless we can see a particular object as a whole, on a particular occasion.

This provides the basis for the final conclusion, about *when* we see certain kinds of things. Seeing something does not entail seeing or having seen how it looks. The case is the reverse: seeing how it looks entails seeing it, and subjecting it to a certain kind of inspection. This inspection demonstrates further visual skill than what it takes simply to see the thing.

Seeing something entails the *potential* to see how it looks, however. And when we exploit this potential, developing our visual capacity into an ability to see more, there will be something to say about the way that development comes about, as well as a way to see how the visible world affords that development. Visual exploration will never come to a

[15] For an illuminating elaboration of these ideas about seeing the shapes of opaque things, see Siewert (2006).

conclusion; and the visible *world* will always have been seen only partially. But a clear view of an opaque object before us is securely behind us.

Chapter 3
Seeing Things in Pictures, and Pictures in Things

Fig. 1. Xu Beihong, *Galloping Horse*. 1941

In the previous chapter, I probed a notable tendency in philosophical discussions of perception to say that what you can see is nothing other than the surface of the things in front of you. We saw that many philosophers have wanted to reject this claim, and to maintain instead that you can see an opaque object, like a tomato, when it sits in front of you: not merely a surface, but the object that is the bearer of that surface. By investigating this dispute, I tried to clarify the relationship between two kinds of experiences: seeing things, and seeing how they look. The investigation made clear that the latter depends on more sophisticated visual skill than the former; it is the manifestation, or the reward, of further development of our capacity for sight. In this chapter, I will continue to seek clarification of the kind of development that our visual capacity admits of, and of the phenomenon of seeing how things look. I will do this by turning to a different, less discussed challenge to the idea that all we ever see are two-dimensional expanses. This one arises through reflection on a certain kind of visual experience, one that occurs in the context of looking at a particular kind of visible object, namely, a picture. In thinking about what this experience is like, we must eventually recognize it as a paradigm case of seeing how something looks. Not only does this give us a way to think about what it is to see how something looks, it also serves to illustrate the phenomenon of being *shown* something, and the business of *making* things visually salient. In this way, reflecting on pictorial representation and our perception of pictures helps us understand how our capacity for sight can develop and how come to see the things we do.

I can introduce the issue that will get us there by posing a seemingly simple question: When you look at a picture, what can you see?[16] Imagine the picture is the one whose image is reproduced above, made by the Chinese artist Xu Beihong in 1941, with ink and watercolor on paper. What you are looking at is a thin, opaque object, whose broad, flat surface supports a configuration of marks made with pigment. Let's imagine that it's right in front of you, and illuminated well.

To say that in this scenario the surface of the object in front of you is *all you can see* raises suspicion, but not (or not only) because it seems you can see the object whose surface it is – the piece of paper itself. Pictures offer us an experience that seems notably enriched or complex in comparison to that of seeing a mere two-dimensional expanse. But in the case of looking at a picture, the inclination to say that we can see something "beyond" its surface is quite different from the inclination to say that we see it as belonging to a voluminous whole. Pictures are creatures of the second dimension, in that they are found on the surfaces of things, and specifically on parts of them that are relatively broad and flat. In order to have a picture in view, seeing the object that it appears on is not enough. One must be facing the part of the object's surface on which the picture is composed. Seeing a picture involves the experience of seeing a part of something's surface. Yet it is evidently distinct from merely seeing a two-dimensional expanse. What more it involves is not a matter of what lies beneath or behind that surface, or the object that it belongs to. Instead it has to do with what is *on* the surface – what the surface itself is like. It has to do with the fact that the surface is demarcated. But a picture provides more than simply the experience of seeing markings on a surface. There are plenty of objects with surfaces, including many with marked surfaces, that do not depict anything, and the perceptual experience they afford seems to lack something that a picture like Xu's painting offers.

Yet it proves difficult to say just how the latter experience is supplemented in comparison to the other – to say in what sense, if any, pictures provide the opportunity to see more than what lies on their surfaces. Wittgenstein expresses the frustration that one can feel in thinking about this question in a remark in *Philosophy of Psychology*.[17] He asks:

> When I see the picture of the galloping horse – do I only *know* that this is the kind of movement meant? Is it superstition to think I *see* the horse galloping in the picture? – and does my visual impression gallop too? (Wittgenstein 2009, 175)

The first question suggests one way of accounting for the intuition that seeing something as depicting what it does is enriched in some particular way. It is not just a matter of seeing the picture, but also of knowing something about it: what it "means," or depicts. We might think

[16] This set up bears some similarities to Noë's introduction of the topic of depiction in "Presence in Pictures," which begins with the question "What do you see when you look at a picture?" (Noë 2012, 83). Here, I ask what you *can* see, rather than what you *do* see, because it is important to my approach that when you look at a visible thing you have the *opportunity* to see it, but there is no guarantee that you do. (This is not something Noë denies, but it is not so much of a concern in his discussion.) Noë's discussion of this question proceeds quite differently from the one in this chapter, in that his initial example is of a photograph of Hillary Clinton – a picture of a well-known person – and he focuses in large part on the relation one has to her in virtue of looking at her picture (he explains it as a form of "access"). This chapter focuses on seeing a picture simply as *depicting something*, not as portraying a particular individual. The difference between these two phenomena will be discussed in more depth in the following chapters. For now, I hope that the example I give will be sufficient to pick out the phenomenon of interest.

[17] In the Wiley-Blackwell 2009 edition of *Philosophical Investigations* this section of the text is titled "Philosophy of Psychology"; previous editions referred to as "Philosophical Investigations Part II."

of the seeing and the knowledge as intimately related: perhaps seeing the picture puts one in a position to know what it means, and perhaps the knowledge colors, so to speak, the experience of seeing. As to how seeing the picture might put one in a position to know what it depicts, there are various ways we might understand this. Seeing the picture might provide one with evidence, on the basis of which one can deduce or somehow figure out what it depicts. Or seeing it might enable one to learn what it depicts in the way that one apprehends things by reading.[18] But however we conceive of it, on this view what is known goes beyond what is seen. Only the marked surface is seen; what it depicts is not seen, but merely known.

Wittgenstein's second question might seem to encourage this view of the experience in question. It mentions an alternative, and suggests immediately that to accept it would be to succumb to magical thinking of some kind. The alternative is the horse that one can see the horse galloping in the picture. On this way of thinking, the picture's object is visible to someone looking at the picture. Wittgenstein's question gives voice to a feeling that no respectable thinker could accept this. He doesn't say what kind of magic someone who accepts it would be thinking of. But one can imagine candidates: the belief that a combination of pigment and paper, by some alchemy, can conjure a horse. Or, their combination summons a horse from somewhere else by a paranormal process. Or, that they induce a "vision" in those whose eyes fall on them, by some equally paranormal process. It would certainly be superstitious to think any of those things are true.

But it is not certain that these are the only ways to interpret the statement that a person who looks at a picture of a galloping horse can see the horse galloping in the picture. It does not rule out that the best conception of our experience of pictures has it that depicted objects are not merely known about, but are themselves visible. In fact, the investigation that follows this remark in *Philosophy of Psychology* does not suggest that this kind of conception should be dismissed so quickly. Wittgenstein goes on to distinguish subtly different senses of the verb "see," corresponding to different kinds of objects that we can be said to see, or different ways in which we can be said to see things.[19] This is precisely the kind of work that needs to be done to determine whether it makes sense to say that depicted objects are visible.

However, the work that Wittgenstein does in these passages does not take us all the way to an answer to that question. His overarching interest is in visual phenomena that do not essentially involve pictures: seeing a resemblance between things, seeing aspects of things, and seeing something as a thing of a certain type. He uses pictures as examples, to illustrate these phenomena and clarify our ways of talking about them, but he does not go on to consider explicitly what it is to see something as *a picture*, as opposed to a different kind of thing. That is not to deny that his discussion of seeing-as provides resources for thinking about the visual experience of depiction; it would be fruitful, I think, to examine it in detail and clarify its relevance for our experience of seeing pictures. It would be particularly interesting to ask why Wittgenstein uses pictures as examples – whether and why they illustrate the phenomenon especially easily, or vividly. But this is not the route I will take in this chapter. Instead, I want to focus on the work of Richard Wollheim, a philosopher who is known for championing the idea that the experience of seeing a marked surface as

[18] The next chapter will deal with the difference between pictures and words in more detail.

[19] It is important that these are two different ways of putting the same point. The distinctions that Wittgenstein makes between different kinds of visible objects are *also* distinctions between different kinds of seeing, or senses of "see."

depicting what it does is enriched by seeing the depicted object as distinct from the markings that depict it. His views on pictorial perception have been hugely influential, and he is clearly attuned to some compelling intuitions. But his views have also met with widespread dissatisfaction. By examining his views and the objections they have faced, I'll try to distinguish insights from missteps when it comes to thinking about what it is like to look at pictures. In doing this, I aim to accomplish three things. One, I hope to further understanding of the nature of the experience of seeing surfaces, the varieties it comes in, and its place in our visual lives. Two, I aim to do the same for the experience of seeing how things look. Three, I hope to further understanding of Wollheim's work, by presenting a way of interpreting and assessing Wollheim's view of depiction that differs from the standard one.

I'll do this in the three sections that follow. In the first, I'll present Wollheim's subtle and systematic way of thinking about the experience that pictures afford. I'll emphasize the framework he introduces for classifying experiences according to their phenomenology, and the way he fills it out by reference to the perceptual capacity that underwrites them. I'll bring out the importance of these aspects of his discussion in section 2, when I turn to Malcolm Budd's objections. Budd's response to Wollheim's views are representative of a common form of opposition to them, and I'll argue that this opposition actually mislocates the problem. Understanding Wollheim's framework resolves the issues Budd raises, and reveals that Wollheim's characterization of the perceptual capacity that underlies our experience is the real stumbling block. In section 3, I'll present an alternative way of characterizing the experience of seeing something as depicting what it does, by bringing to Wollheim's general strategy some useful concepts that he overlooks – namely those of *images* and *display*.

1. Wollheim's framework

Wollheim is led to the questions of perceptual phenomenology that I have raised in the interest of coming up with an analysis of what a picture is. Starting with the definition of a picture as an object that represents something in a particular way, Wollheim claims that the kind of representation that makes something a picture is "to be understood through, though not exclusively through, a certain species of seeing" (Wollheim 1980, 205). The species of seeing he has in mind is precisely the kind of experience we have been considering: seeing something as depicting what it does. [20] The occurrence of this kind of experience is something that a picture makes possible, but not something that it guarantees. Something could be a picture without ever being seen; equally, those who see a picture might fail to perceive it as depicting anything, or they might misperceive what it depicts. So what Wollheim is suggesting is that depiction is to be explained as a kind of potential that things can have – the potential for experience with a certain kind of phenomenology.

Wollheim's strategy for pinpointing the kind of experience in question involves relating it to a broader class to which it belongs: a genus of which it is one species among others. The genus is a perceptual one; that is, everything that belongs to it is a form of

[20] The grammar of the phrase "seeing something as depicting what it does" may be somewhat cumbersome, but it is very useful in this context. For one thing, it implies that what is seen *is* a picture, not just that it looks like one. And it implies that the picture is seen not only *as* a picture, but as the picture it is, so to speak. It is seen as depicting what it in fact depicts. The fact that pictures can be seen "incorrectly," as depicting things other than what they depict, has led philosophers to put forth various claims about what determines the content of a picture. Some (like Wollheim) say that the artist's intention is a deciding factor, others point to features of the picture's context. Using this locution we can sidestep the issue of choosing between these alternatives.

seeing, or visual perception. But it is one genus among others, a narrower category than is picked out by the idea of seeing in general, or as such. Thus Wollheim's phenomenological work crucially involves the construction of a framework for classifying experiences. And the framework is constructed by way of comparisons: by picking out examples of other kinds of visual experience, and articulating their similarities and differences in a way that identifies their theoretically significant characteristics.

Ultimately, Wollheim classifies seeing things in pictures as a species of a *perceptual* genus. That is, everything that belongs to it is a form of seeing, or visual perception, a sub-category of the "family" that contains seeing in general, or as such. To make his case, he starts by referring us to various forms of experience that count as species within the genus of interest. One is "the seeing appropriate to photographs," or seeing photographs "as photographs." Another is "the perception of Rorschach tests." What Wollheim has in mind here is the *taking* of Rorschach tests – which involves looking at cards printed with "ink blots"[21] and responding to the question "What might this be?" A final example is an activity that Leonardo da Vinci recommends in his *Trattato*, a manual on painting: gazing at a "damp-stained wall" or "stones of broken color" and "discerning there" things like "scenes of battle or violent action and mysterious landscapes" (Wollheim 1980, 218). This is supposed to provide inspiration for new motifs.

However difficult it may turn out to be to state what unifies the phenomenon that these examples illustrate, their presentation is meant to forge our initial grip on the genus to which they all allegedly belong. Characterizing what distinguishes this genus is one of two ways Wollheim tries to shed light on the kind of experience pictures provide, and make us understand the sense of saying that we can "see" depicted objects in looking at the pictures that depict them. The other way he does this is by articulating what differentiates the pictorial species of the genus from the others.

In the interest of stating explicitly what these cases all seem to involve, Wollheim coins two terms that have since become central to philosophical discussion of depiction. First, he calls the experience that pictures characteristically afford the experience of "seeing-in." This is initially introduced in contrast to "seeing-as," which Wollheim had previously used to characterize the phenomenon by saying that it is "a matter of seeing x (= the medium or representation) as y (= the object, or what is represented)." He had hoped that depiction could be elucidated in this way, since "through the initiative of Wittgenstein, we seemed to be gaining a good understanding" of seeing-as (Wollheim 1980, 209). But he says he has come to realize that this is not a viable analysis of depiction.

The name "seeing-in" derives from the fact that the experience can be described as that of *seeing* one thing *in* another. But the emphasis on this locution has the potential to mislead. Seeing-in differs importantly from other experiences which we might describe as seeing one thing in another, such as seeing one's reflection in a mirror or seeing the sun rise in the morning. It is different in that it involves seeing *both* of the two things referred to, and seeing the one by means of seeing the other. That is, it is a matter of seeing one thing, and also seeing something else in it. In the case of a picture, the something else is a depicted object, or what a particular picture is of.

Wollheim's second coinage is an attempt to mark what is distinctive of the phenomenology of this experience, which makes it different from the others that are

[21] The cards have not actually had ink blotted on them; they are reproductions of shapes that the psychologist Hermann Rorschach made and selected to form a standard collection on the basis of experiment with the examination of schizophrenic patients.

referred to by the same locution. He calls the feature that of being "twofold," or exhibiting "twofoldness" (Wollheim 1980, 219). Part of what he means by this is that the experience is one of seeing *two things*. But it is not meant to be just any such experience – seeing a *pair* of things would not count as having a "twofold phenomenology" as Wollheim intends the term. Twofoldness amounts to having a particular kind of structure in which the two objects are implicated, and in which they play different roles. Twofold visual experience is a matter of seeing one of the things in one's presence, and also seeing something else, whose visibility is "sustained" by seeing the first.

So the things seen on an occasion of "seeing-in" must not only be two in number, they must also be related by playing complementary roles in the experience. Though they are distinct, seeing the second must come along with seeing the first.[22] It is not obvious, however, how to square his claim with the fact that it certainly *is* possible to see a picture without seeing what it depicts (one might be too far away to see enough sufficient detail on its surface, or one might simply fail to perceive its markings as composing a depiction of a particular object, or one might misperceive it as depicting something else). But I think it can be done, and that trying to reconcile them makes clearer what Wollheim is trying to say about the experience of pictures. I think what he wants to say is that seeing a picture as depicting what it does (what we are now considering as a case of "seeing-in") involves seeing a marked surface, but seeing it in a particular way – one which necessarily implies that one also sees something else in it. It is incompatible with seeing the marked surface in what we might call a "straightforward" way, without the involvement of any other object of sight. So although it is possible to see what is in fact a picture in this "straightforward" way, without seeing what it depicts, seeing it *as depicting* does not include that "straightforward" experience as an element. It is a different way of seeing the marked surface – one with a different phenomenological structure. Thus the facts about perceiving pictures without perceiving what they depict are compatible with the point that in general, the way one sees a marked surface, in the case of seeing-in, necessarily involves seeing something else.

All of these references to seeing "something else" might feel a bit vague. One might wonder whether there is a way of picking out what is seen in this way as a certain kind of object. It is not necessarily a *depicted* object, since the genus of seeing-in includes non-pictorial species. So one might expect Wollheim to use a term that subsumes depicted objects under a more general heading. Wollheim does not find or coin such a term, however. He frequently makes use of variables: seeing-in is described as "seeing x in y." He also labels the aspects of the experience of seeing-in that correspond to these of variables: x corresponds to the "recognitional fold," while y corresponds to the "configurational fold."

But the terms "twofoldness," "seeing-in," "recognitional fold," and "configurational fold" simply give us names for the phenomena we want to understand, and they are essentially placeholders, like "x" and "y,". The work of illuminating the phenomenon of pictorial perception has been done not by the coining of these terms, but, so far, by its comparison to certain non-pictorial cases, both within a "genus" and without. Wollheim has suggested that seeing something as depicting what it does is an instance of a more general

[22] In his initial presentation of seeing-in, Wollheim talks of seeing the one thing and seeing the other as two separate experiences, but in *Painting as an Art* he clarifies that they are two aspects of a single experience, which are "distinguishable but also inseparable" (Wollheim 1987, 46). Hyman points out that without an established way of counting experiences, it is not obvious what this revision amounts to (Hyman 2003, 678). I am inclined to think that it does not represent any deep change in Wollheim's thinking, but rather an attempt to say more clearly what he had in mind initially.

phenomenon, an experience with a certain kind of structure. A depicted object is one of the things that can have a place in this structure, but it is not the only one. It is part of a broader category of objects that correspond to a particular sense of the term "see."

Simply recalling the experience of looking at inkblots and wall stains may induce the sense that there is some category to which they all belong. But articulating the bounds of that category that unify the things in it requires further resources. Wollheim attempts to do this by characterizing the experiences in this genus as exercises of a particular fundamental perceptual capacity. The relevant capacity contrasts with what Wollheim calls "straightforward perception," which is "the capacity that we humans and other animals have of perceiving things that are present to the senses" (Wollheim 1980, 217).[23] Wollheim thinks that "straightforward perception" does not exhaust what perception is for us, and that we are also endowed with a "special perceptual capacity" which "allows us to see things not present to the senses" (Wollheim 1980, 217). Seeing-in is the manifestation that capacity. With the distinction in perceptual capacities in hand, we have a new way of contrasting the two objects implicated in seeing-in. We have already considered that seeing the one sustains seeing the other – in that way, the one makes the other visible. Now we are told about another difference between them: the one lies in front of us, while the other is not present.

In addition to the contrast with "straightforward perception," Wollheim relates this special perceptual capacity to other visual phenomena. He says: "If we seek the most primitive instances of the perceptual capacity with which seeing-in is connected, a plausible suggestion is that they are to be found in dreams, day-dreams, and hallucinations." But he denies that these experiences are in fact "instances" of seeing-in, even though they may "anticipate" it or be "continuous with it" (Wollheim 1980, 217). These comments are brief and less precise than one might hope, but Wollheim seems to be acknowledging that there may be some important relationship between the capacity for seeing-in and the imagination, while denying that seeing-in is itself an experience of imagining. The crucial difference, he says, is that hallucination and day dreams "arise simply in the mind's eye," where as seeing-in "come[s] about through looking at things present" (Wollheim 1980, 218). Ultimately, these remarks serve to reinforce the point that seeing-in is a variety of *perceptual* experience, and specifically a kind of *seeing*. In dreaming and hallucinating we do not see, but rather merely *seem to see,* and in day-dreaming we may *imagine* seeing. But seeing-in, Wollheim is saying, is a special case of seeing, and it is the manifestation of a special *perceptual* faculty, rather than an

[23] Wollheim adds a wager about the best way to understand its nature: "Any single exercise of this capacity is probably best explained in terms of the occurrence of an appropriate perceptual experience and the correct causal link between the experience and the thing or things perceived" (Wollheim 1980, 217). Wollheim's attraction to the causal theory of perception no doubt influences how he thinks and writes about the experience of seeing a picture as depicting what it does. I suspect that it is what leads him to say, occasionally, that pictures are, by definition, designed to *produce* a certain kind of experience *in viewers*. This leads Hyman to read Wollheim as advocating a "subjectivist" view of depiction, according to which depiction is defined in terms of a psychological state that can be triggered by coming into contact with a marked surface (Hyman 2000 and Hyman 2003). Hyman then objects to Wollheim on the grounds that subjectivism is mistaken. But Wollheim does not explicitly commit to or rely on a causal theory of perception in his discussion of depiction. His attraction to it is distant enough from his philosophical interest in depiction that it is fruitful to interpret and evaluate his claims as compatible with the kind of alternative to the causal theory that I consider in chapter 1. And if we draw on that view in understanding what Wollheim means when he defines depiction in terms of a certain kind of visual experience, his view will not seem to be "subjectivist" after all. A picture will be something that is designed to afford the opportunity for a certain kind of experience, but that experience is not a merely subjective state that is merely triggered by contact with an external object.

exercise of the imagination. That faculty is for seeing things that are not present to our senses, in virtue of seeing things that are.

To repeat, pictures are among the things that engage this capacity. Wollheim distinguishes them from others as objects that afford the experience of a particular "species" of seeing-in. It is distinct from other such species in that it is "subject to a standard of correctness" and one that is "set by an intention" (Wollheim 1980, 207). He means that when one sees what is *depicted* in a marked surface, one sees what one is *supposed* to see in the marked surface, that is, one sees what the person who made the markings intended viewers to see. Wollheim does not make this explicit, but we can infer that the standard of correctness makes a difference to the experience from the perspective of the viewer: the viewer *takes* whatever they see to be what they are supposed to see, given the intentions of the person who marked the surface. In the case of seeing-in in the context of a Rorschach test, there is no standard of correctness in play – there is nothing one is supposed to see in the inkblots that are presented. In the case of seeing a photograph, there is a standard of correctness, but according to Wollheim its basis is different: it is "in large part a matter of who or what engaged in the right way with the causal processes realized by the camera" (Wollheim 1980, 208). Thus the species of seeing-in are phenomenologically different, though they all share a fundamental phenomenological structure, as exercises of the same fundamental perceptual capacity. By grasping what unites them as well as how they are distinct, Wollheim thinks, we can understand the nature of the experience that pictures characteristically afford. And he thinks this is also an understanding of what depiction is.

If Wollheim's views were correct, they would give us a way of maintaining that we can see the horse galloping in Xu's picture, assuaging the worry that such a statement can only express superstition. Many of Wollheim's readers, however, have not found his explication of pictorial experience entirely satisfying, and in some cases its perceived shortcomings have reinforced the appeal of the "seeing + knowing" view of the experience in question. Many have focused on his explication of "twofoldness," and found him to be unjustifiedly "quietist" about what it amounts to. I think these readers are right that what Wollheim says does not, as a whole, constitute a tenable alternative. But I think that these readers have rejected Wollheim's views in a way that largely misconstrues his arguments and overlooks his insights, and in a way that effectively throws the baby out with the bathwater. The shortcomings of Wollheim's account, though they are there to be found, should not make us give up on the idea that there is a special but intelligible sense in which depicted objects are visible, and that pictorial perception can involve seeing – really seeing – a depicted object.

To make my case, I'll examine Malcolm Budd's criticism of Wollheim, which is representative of a popular way of disputing Wollheim's view. Identifying problems in Budd's objections will put us on track to find a better understanding of a depicted object as a possible object of sight, or equivalently, a better understanding of how such a thing can be seen.

2. Seeing-in and awareness of appearance

In "On Looking at a Picture," Budd asks the question that I began with: what can you see when you look at a picture – and more specifically, what is the phenomenology of seeing something as depicting what it does? He voices at the outset the intuition that the experience of seeing something as depicting what it does is enriched, in some way, in comparison to that of seeing something like a "mere configuration of lines" (something which "means nothing to you") (Budd 2009, 185). According to him, this enrichment can be

uncontroversially referred to as a matter of "visual awareness" of something other than a marked surface. Specifically, it involves visual awareness of the look of something – the look of what is depicted. Taking as an example Claude Monet's *The Seine in Thaw*, he says, the experience of seeing it as depicting what it does must "involve both an awareness of the marked surface and an awareness of what a thawing river looks like" (Budd 2008, 185). But for him, this does not yet constitute an answer to the question, because of an ambiguity that attaches to the idea of visual awareness. The challenge of understanding experience of this kind is, then, to identify which kind of visual awareness is at issue, and so to "explain in what way your experience of Monet's picture involves a visual awareness of a marked surface and also a visual awareness of a river in thaw" (Budd 2008, 185).

> As Budd understands the term "visual awareness," it
> …covers both experiential and dispositional forms of visual awareness – experiential, as when you see a river or visualize one in the mind's eye; dispositional, as when you possess the capacity to recognize a river if you see one or to recall in your mind's eye how a river looks. (Budd 2008, 186)

So the question arises: how are each of these phenomena involved experience with pictures? Are we "experientially" or "dispositionally" aware of the marked surface of the picture? And what about whatever it depicts?

Budd reads Wollheim as answering this question by saying that a viewer is experientially aware of both a marked surface and a depicted object in the experience at issue. This is appropriate insofar as Wollheim's hierarchy of perceptual species and genera is meant to classify forms of "occurrent" experience, rather than recognitional dispositions or forms of implicit familiarity. But I think that to read Wollheim's discussion of seeing-in as an answer to the question of which of these two forms of visual awareness is at issue is effectively to ignore the phenomenological framework that discussion presents. In what follows, I'll explain why.

Budd comments on Wollheim's earlier and later discussions separately. The difference between them has been mentioned already: in later work, Wollheim shifts from describing seeing-in as two distinct but simultaneous experiences of seeing to describing it as a single occurrence with two distinct aspects. When it comes to both versions, Budd judges Wollheim to be problematically quietist, leaving things in need of development which they cannot be given.

When it comes to the first formulation, Budd objects that there is a "lacuna in the account" Wollheim gives, in that the nature of the experience of the depicted object "has been left blank, and it is difficult to see how it could possibly be filled in" (Budd 2008, 196). What's worse, it is "difficult to see how it could be filled in" (Budd 2008, 196). That is, not enough is said about *how* the depicted object appears, or what it is like to see it, and it seems there is simply nothing helpful that *could* be said to this end. Budd supports this by considering two ways of trying to accomplish the filling in required, and explaining why they are unlikely to succeed. One is to say that it is "indistinguishable by the subject from a corresponding instance of face-to-face seeing." Budd's thought is that an instance of face-to-face seeing that corresponds to seeing a picture of, say, a horse as depicting what it does is an experience of seeing a horse. The experience couldn't be that (a picture of a horse is not a horse[24]). But it could instead be the experience of *seeming* to see a horse and being unable to

[24] Cf Magritte, *The Treachery of Images*.

tell that one is not. But as Budd notes, this would be "inconsistent with Wollheim's longstanding opposition to illusionistic accounts" (Budd 2008, 196).[25]

The other option Budd considers is that the experience is that of *visualizing* something. He then notes phenomenological differences between visualizing something and seeing something that depicts it: "your capacity to visualize a complex state of affairs is likely to be extremely limited, and in so far as you succeed in visualizing the state of affairs your image is likely to be infected by a considerable degree of vagueness." In the case of seeing a picture of a complex state of affairs, on the other hand, you "normally have little if any difficulty" seeing it as depicting what it does in its full complexity (Budd 2008, 197). So the awareness of the object depicted cannot be characterized as a matter of visual illusion, nor as an episode of visualization. In the absence of another way of saying what kind of experience it is, the claim that there is such an experience cannot be defended. But no answer is forthcoming, according to Budd.

This objection is unsatisfying, however, because it seems to overlook much of the work that Wollheim does to characterize seeing-in. Given how Wollheim has presented it as a "genus," the question of the aptness of either of these phenomenological proposals (assimilation to illusion or visualization) should not arise. Indeed, Wolheim does not just reject the conception of pictorial experience as illusory, his view is explicitly presented as an alternative.

As I explained in the previous section, Wollheim thinks that the experience of seeing something in a picture is a member of a certain *perceptual* genus – a certain kind of *seeing*. Illusion (seeming to see) and visualization (imagining seeing) are categories of visual experience that each contrast with the genuinely perceptual (seeing). The conceit of Wollheim's framework is that the category of seeing itself admits of philosophically significant divisions (and that within those there are further ones as well). Wollheim's division of the category of visual perception into significant classes invokes an idea that Budd does not consider in his assessment, though he reports it a few paragraphs earlier: the idea of "visions of things not present [which] come about through looking at things present" (Wollheim 1980, 218). By "visions," Wollheim does not mean hallucinations, illusions, or visualizations. He has in mind the genuine perception – seeing – of things which are not present.

This is, at least, a grammatically well-formed way of saying what the experience is, as distinct from visual illusion, visualization, and seeing things that are present. If it does not actually describe a form that perception can take, something must be said to show why not. To that end, one could object to the idea that seeing admits of different species, by arguing that it has the same phenomenological structure in every case. But if it does make sense to distinguish species of seeing, it is illegitimate to insist that they be distinguished in terms of the concepts of illusion and visualization, since these concepts categorize visual experience at a more general level. So pointing out that seeing-in, as Wollheim discusses it, does not fit the description of these other kinds of visual experience could not constitute a legitimate objection to Wollheim's view. Perhaps we should hope for a clearer explanation of seeing-in as distinct from straightforward seeing, or a fuller argument that the experiences in question

[25] The basic problem with this view, and the gist of Wollheim's objections to it, is that considered in this way, having the experience in question would be incompatible with seeing the marked surface for what it is, and (thus) incompatible with having the experience of seeing something as depicting what it does. (See Wollheim 2003 for a more extended discussion.) Gombrich 1960 is the usual target of criticisms of the "illusion theory" of depiction; for a discussion of whether Gombrich is properly interpreted by these critics, see Bantinaki 2007.

form a unified category. But we should not demand an answer to the question of whether seeing-in is supposed to be a form of illusion or a form of visualization.

Budd's criticism of Wollheim's later discussion of the recognitional and configurational aspects of the experience also demands development in the wrong direction. Wollheim distinguishes these aspects by saying that the one can be understood through an analogy with seeing a meaningless marked surface, the other through an analogy with seeing "face-to-face" something that fits the description of the object in the marks. But then he cautions that these respective aspects and these comparison cases are actually "incommensurate," since the aspects must co-occur, while the comparison cases can occur separately. On that basis, Budd explains, Wollheim regards it "illegitimate to enquire about the experiential resemblance between either aspect of the complex experience and the simple face-to-face experience after which it is described" (Budd 2008, 200).

Budd then objects that this is not a good reason for rejecting questions about the "experiential resemblance between the recognitional aspect (for example) of a seeing-in experience and the face-to-face experience after which it is described" (Budd 2008, 200). He protests: "For whatever the particular complexity of the twofold experience is supposed to be, why should it render the phenomenology *of its recognitional aspect* incommensurate with that of the face-to-face experience?" (Budd 2008, 200). We are left without an answer to a question we need answered, if we are to grasp the nature of either aspect, and through that the phenomenology of the experience as a whole. Thus the problem of "blankness" is supposed to attend this articulation as well.

It must be acknowledged that Wollheim makes himself less clear on these points than one might have hoped he would. However, I think this objection misconstrues how the analogies in question should work. Here is what Wollheim says about the matter in *Painting as an Art*:

> The two things that happen when I look at, for instance, the stained wall are, it must be said, two aspects of a *single experience* that I have, and the two aspects are distinguishable but also inseparable…They are neither two separate simultaneous experiences, which I somehow hold in the mind at once, nor two separate alternating experiences, between which I oscillate – though *it is true that each aspect of the single experience can be described as though it were a case of simply looking at a wall or a case of seeing a boy face-to-face*. But it is an error to think that this is what it is. (Wollheim 1986, 46)

What does it mean to say that an aspect of an experience "could be described as though it were a case" of a different kind of experience? What does it mean to add to this statement that "it is an error to think that this is what it is"? It seems that Wollheim can only be denying that the aspects of the complex experience at issue are, or are relevantly like, the other experiences that have been mentioned. But to deny this is perfectly legitimate – it just doesn't get us very far. It's legitimate for the following reason: if understanding the complexity of this phenomenon is a matter of prising apart two aspects of a unified experience, then we should not think of or describe those aspects as we would think of or describe experiences in their own right. It is a mistake – a category mistake – to put them on a par with cases of experience and compare them. The claim that the "simple face-to-face experiences" mentioned are "incommensurate" with the aspects of seeing-in is, in the end, a rather roundabout way of reiterating the claim that what we are interested in is itself an experience, rather than a collection or sequence of experiences. In that sense, these

suggested analogies don't get us any further in understanding what the relevant experience is like.

It might seem like they could, despite the categorical difference between experiences and their aspects. After all, it would be perfectly legitimate to compare the experience of "seeing-in" *as a whole* with, on the one hand, looking at a set of meaningless marks, and on the other, seeing something face-to-face. And one could say that making each of these comparisons allows us to identify some salient similarity, which corresponds to or puts us in mind of either of the two significant aspects which give this experience its distinctive phenomenological structure. Wollheim did not say this, at least not clearly – but it is compatible with what he did clearly commit to.

However, these comparisons simply can't be of much use. Both of them would seem to count as cases of what Wollheim has called "straightforward perception," and it is the *contrast* with this form of experience that we must grasp if Wollheim's framework is to elucidate things for us. The comparison between cases of these two kinds is where we have to begin, and what we must do is identify the theoretically relevant *differences*, not say how they are analogous. To identify their similarities would only be a distraction from the phenomenological tools that Wollheim has already provided: the distinctions between the various "species" of seeing-in, and the idea of seeing-in as one "genus" of visual perception among others. Picking out the genus requires *contrasting* seeing things in marked surfaces with seeing marked surfaces without seeing anything in them. And picking out the species requires further contrast, between seeing things in *pictures* and seeing things in non-pictorial marked surfaces.

Budd's objections, and similar complaints that Wollheim is unacceptably quietist, ultimately fail to bring into focus both the merit and the crucial misstep in Wollheim's account. They miss the mark in large part because of a reluctance to accept and address the fact that Wollheim is committed to the claim that when one sees something as depicting what it does, one visually perceives a depicted object. This is what leads Budd to search in vain for some other form of "experiential awareness" that Wollheim could invoke to characterize the relationship with depicted objects. But on Wollheim's view, the relationship is that of *seeing* – in a perfectly literal, though special, sense of the word "see." Wollheim's claims cannot be evaluated without acknowledging this; this commitment is at the heart of the account. Yet Wollheim's point in identifying seeing-in as a "perceptual genus" goes unrecognized by these readers, and so they do not properly appreciate the promise of his strategy for accomplishing the phenomenological work of characterizing seeing something as depicting what it does.

At the same time, this oversight distracts from the major problem with Wollheim's view. Only when his strategy is properly understood does his way of distinguishing seeing-in from straightforward perception become clear, and clearly false. He claims that seeing-in is not only an experience of seeing a marked surface, it is also an experience of seeing something not present to the senses.[26] And the idea of this kind of vision constitutes an unacceptable strand in Wollheim's thinking, more or less a form of magical thinking. It is different from the forms of superstition about pictures that I outlined earlier: thinking that there are ways of combining things like pen and paper to forge instances of things like

[26] To be precise: it is not that it contains these two as elements, it simply is itself both – it is *the experience* of seeing *two such things*, one made visible by seeing the other. The claim also serves to draw the distinction between the configurational and recognitional aspects of the experience: the former amounts the status of seeing of a marked surface located nearby, the latter to the status of seeing of something not present.

horses and mountains, or to summon them from elsewhere. Here, the idea seems to be that the things pictures depict are there to see, when we look at those pictures, but they are not in our presence – they are visible, but they occupy some kind of distinct realm that is discontinuous with our surroundings. But this idea is objectionable in much the same way as the others: it effectively assimilates the seeing of depicted objects to the seeing of things in crystal balls. Accepting it requires us that we depart from common sense, and give up the conceptual truth that you see only what is present to your senses; to see something just is for it to be present to your visual sense.[27]

Budd's reluctance to acknowledge straightforwardly that Wollheim is committed to regarding depicted objects as visible may derive from the fact that he and Wollheim share the supposition that whatever a picture depicts, it is not *there*, when and where the picture is. On the basis of this commitment, Budd may see no distance between admitting that the depicted object is not there, and recognizing that it cannot be seen. He is right to insist that there is none – but mistaken in assuming that Wollheim agrees.

But this flaw in Wollheim's view does not show that he is wrong about whether we can see depicted objects when we look at pictures, and that Budd is right that the "experiential awareness" we enjoy in looking at them excludes them, extending to nothing more than the marked surfaces that depict them. It does not speak against the view that there is an important difference between two structures of visual perceptual experience (two "perceptual genuses"), which can help us to understand what happens when we see something as depicting what it does. Nor does it rule out that the difference is made by the visibility of something in addition to a marked surface – that the experience of seeing something as depicting what it does entails seeing two things, one of them a depicted object. It simply means that the matter rests on whether we can make sense of depicted objects as *being there*, when a picture is present.

Wollheim gives us a strategy for making sense of this: subsuming the way in which depicted objects can be visually present under a more general, but still distinct, category of perceptual presence. From there, the way forward does not lie in coining more technical terminology, nor in the acceptance of mysterious forms of visual experience. Rather, it lies in considering the variety of examples identified as instances of the general category "seeing-in," and looking for the right *familiar* terms to express what is characteristic of them. Wollheim himself does not use this strategy to an entirely satisfactory effect. His claim that depicted objects are things seen when not present to the senses does not constitute a tenable conception of what a depicted object is as an object of sight. To understand what seeing-in is, we need an adequate conception of what special kind of visible object is involved in seeing-in.

In the next section I will try to do better, by sketching a way of making sense of how a depicted object *is* there, in virtue of a picture's presence. My suggestion will draw on the familiar notion of an image, as well as some general reflections on visual appearances. But before turning to the positive proposal I want to present briefly some considerations that might have stopped Wollheim from going this route himself. I think that they touch on legitimate worries, and I want to make clear how these worries can be addressed.

It is possible that Wollheim does not pursue the issue in this way because he is wary of circumscribing too narrowly the range of things that pictures can be of. Pictures can depict things of unfathomably many different types, from horses to horse races to unicorns

[27] To say this is not to endorse a "causal theory" of perception, or any form of perception. "Present to" does not pick out a causal relation. For a discussion of the causal theory of perception, see Chapter 1.

to uprisings. To say *what kind of thing a depicted object is* might seem to rule out possibilities for depiction from the armchair, when the limits of depiction are best explored by picture-makers rather than philosophers. But no answer to the question as I understand it would limit the possibilities for depiction to objects of any particular type, or circumscribe the kind of thing a picture can depict in any way. It does not require saying anything about the qualities that a depicted object may or must exhibit, or the concepts that may or must be used to describe what a picture depicts. Moreover, the question does not presuppose that there is anything to say about this. The question to be answered is not about what a certain class of visible objects are *like*, but what they are *as* visible objects, or in other words, what is distinctive about *how* they can be seen.

It is also possible that Wollheim does not pursue the issue because he is committed to articulating a definition of depiction in terms of a certain kind of experience, and he thinks that if the account of this experience requires specifying a certain kind of visible object, he will not have made good on his commitment. But (to repeat) the question does not ask for a specification of the qualities of depicted objects, and so an account that answers it will not result in a definition of depiction that goes beyond the nature of a certain kind of experience. It will still amount to an understanding of depiction in terms of experience with a certain phenomenological structure. The fact that this requires characterizing the role that visible objects play in the experience should not be a problem. The account will be no less "experiential" if the experience in question is not a purely subjective state, that can be specified without reference to anything but the viewer.

3. Seeing-in and images

I now want to return to Wollheim's list of examples of species of seeing-in, to assess what can be said to unify them. They include (though they are not necessarily limited to) cases of looking at photographs and seeing what they are photographs of, taking of Rorschach tests, and the activity mentioned in Leonardo's *Trattato*. Reflecting on these cases, I think we do have the intuition that they are importantly and interestingly similar. They all involve looking at *marked surfaces*, but this doesn't exhaust the similarity we notice in them. They each exemplify a particular *way* of looking at, a particular kind of perceptual experience that some marked surfaces afford. If we follow Wollheim in saying that it involves seeing something further *in* a marked surface, how, in general, does this add to what we see?

Two of these non-pictorial cases of seeing-in seem almost to obstruct any attempt to answer this question. The subject of the Rorschach test is presented with amorphous designs and asked the extremely vague question, "What might this be?" The question is meant to prompt a specific kind of response, and the success of the test seems to depend on the subject grasping what her demonstrative should pick out when she responds with some statement of the form "*That is a*…" But the sense of the question is not explicated by the examiner. And insofar as the test is intended to ascertain something about the subject's idiosyncratic ways of looking at things, the demand that the question makes must leave a fair amount of leeway in what the subject must do to meet it.

The instructions from the *Trattato* are similarly open-ended. The purpose of the activity Leonardo recommends is to find inspiration for new motifs and techniques for painting. The thought behind it is that surfaces embellished naturally or haphazardly may exhibit features that haven't yet been applied to canvases, but which could be – or would be promising starting points for elaboration. Arguably, Leonardo's advice implies that there is a way of looking at non-pictorial surfaces that has some important affinity to the way we see

paintings. But rather than aiming at any precise circumscription of this experience, it emphasizes that there is an important indefiniteness to it, by claiming the value of continually seeking unfamiliar forms of it.

But the case of looking at a photograph is, I think, better poised to point us to the key concepts for relating the "species" as members of the relevant "perceptual genus." This is because we are accustomed to thinking of photographs as *images*. Making use of the concept of an image, and considering what Wollheim calls seeing-in as the experience of seeing images as such, will help to shed light on the phenomenology at issue.

The notion of an image is familiar, but notoriously slippery. The word "image" is used to refer to a collection of heterogeneous phenomena, which are united by a tangled set of conceptual practices. It is important to acknowledge this, and realize that our pretheoretical intuitions about images can only get us so far. Nonetheless, they provide us with a starting point, and careful reflection on our everyday thought about images provides a foothold for clarifying the distinctions we are interested in. In doing so, our use of the term "image" may come to be that of a term of art. But the path we take in getting there should give us a better sense of the significance of various other terms of art, which we will then be in a position to use clearly in thinking about pictures and the experience they afford.

In his book on photography as an image-making technology, Patrick Maynard provides a number of concise and powerful explications of the basic elements of the phenomenon we are interested in. I will make extensive use of his work, but will modify his construals in certain key ways, to better address the concerns at hand.

Maynard begins with an essential clarification of the kind of images at issue, the kind of which photographs and "handmade" pictures are instances. He points out: "they are not "mental images, nor are they mirror images. They are usually physical states of surfaces of which people make mental (that is, cognitive) use" (Maynard 1997, 24). The relevant state of the surfaces in question is that of being *marked*. Marks, in the relevant sense, are "discontinuous physical states of the surfaces, notably but not exclusively those that have appeared due to some action of agencies on the surface" (Maynard 1997, 25). This would include things like "skid, tooth, and water marks but also birth marks." These sorts of discontinuities can be described in terms of texture and chemical composition, and as the products of many different kinds of physical processes. The paradigm cases are relatively permanent alterations to thing's surfaces, which are naturally considered to be parts of the things whose surfaces they lie on. But the category extends to discontinuities of a quite different type, namely cast shadows (such as the "fugitive shadow projections that make up cinematic images on screens") and the physical variations across "emission screens, such as TV" (Maynard 1997, 25). These last examples are discontinuities of an optical nature, and because they are more fleeting, and are "not usually considered to be parts" of the surfaces they are found on, they count as marks only in a somewhat extended sense. But this expansion of the category does not encompass all optical phenomena on surfaces: it excludes reflections and highlights (Maynard 1997, 25).[28]

[28] Maynard does not fully explain the reason for this, and it may seem puzzling or arbitrary. Here is a suggestion about the principle behind the difference: The region of an object on which a shadow is cast can be specified, whereas when an object is reflected in a mirror, there is no particular part of the mirror that reflects it. The location of a cast shadow is determined by the placement of the light source and the object casting the shadow; anyone who wants to see the shadow must look at that part of the surface. In contrast, where in a mirror you must look to see a particular object's image depends on your position, not just the position of the mirror in relation to the object. Similarly, when an object glints in the light, the way in which a view of the object's

Images (of the kind under discussion) are made of marks, but images are not themselves marks, for the following reason:

> We make images by making marks on surfaces, and images are destroyed as these marks are effaced, but it would be a mistake to think of the image simply as a kind of mark. What is the relationship? Clearly, *unmarked* parts of surfaces make up parts of images.[29] Therefore we should think of images, in many cases, as the marked surfaces themselves, or parts of them. (Maynard 1997, 26)

This clarifies the kind of thing an image is to some extent. But it does not give us all we might want, if we think that not all marked surfaces are images. If only a subset of the marked surfaces are images, then we have yet to face the question of what distinguishes those that are images from all the rest. And it seems we don't consider just any marked surface to be an image. This is made clear by the fact that if something is an image, then there is something to be said about what it is an image *of* – and we are at a loss to say anything about what the average stretch of exposed wood grain is *of*. So an image is a special *kind* of marked surface. But what's so special about it?

Another of Maynard's observations about images points us in a promising direction. In light of the reasoning above, about why images should not be equated with marks themselves, Maynard suggests that images are "unities comprising both the marked and unmarked parts of the surface in a single overall appearance" (Maynard 1997, 28). This provides a way of distinguishing them – if what it is to have a "single overall appearance" is something that only some marked surfaces can claim. On the basis of the work of previous chapters on visual salience and its relationship to visual appearance, we can make out a way in which it is distinctive, and how it can account for something like the "twofold" phenomenology of "seeing-in."

I am now in a position to make my suggestion about what is characteristic of seeing images as the images they are. In the first chapter, I showed that a thing's visual appearance can be understood in terms of the circumstances in which it can be seen: its visual salience in those circumstances, relative to the visual skill of potential viewers. In the second, I clarified the difference between seeing a thing and seeing how it looks, or seeing its visual appearance, and showed that the latter is harder, or manifests more sophisticated visual skill, than the former. These points are useful now in the context of understanding the kind of visual opportunity that images characteristically afford. It is an instance of seeing something, and seeing how it looks. It amounts to seeing not just the marked surface on which the image lies, but moreover, seeing how it looks. Its look, or its visual appearance, is what one sees in "seeing-in," in addition to the marked surface.

This is not all that needs to be said about this kind of experience, however. After all, any marked surface, insofar as it is visible as such, has a visual appearance of some kind. And so it is in principle possible to see its look, when one is in a position to see it. In the case of an image, however, one sees something with a particularly *salient* visual appearance: something whose visibility is structured in a particularly (visually) obvious way. The elements of an image – the marks and the unmarked parts of this kind of surface – are *unified* in such a way that how the surface looks is particularly easy to see. The distance between what it takes

---

surface is obscured by highlights depends on vantage point, so in a way there is no particular part of it that is highlighted.

[29] The picture reproduced at the beginning of this chapter serves as an example of this.

to see them and what it takes to see how they look is smaller than it is with other things. In short, with images, it is relatively easy to see how they look.

So, we can characterize the "genus" of perception that our experience of pictures belongs to by identifying it with the perception of images as the kind of images they are – as images of what they are images of, as having the objects that they do. This contrasts with "straightforward" perception in that it amounts to seeing not simply visible objects, but their visual appearances as well. We can also use this understanding of images to characterize the two things whose simultaneous visibility accounts for the fact that what is seen on these occasions is "twofold." We can say that one is the marked surface that constitutes an image, and the other is whatever that image is of. And we can unpack that by saying that the former is a marked surface, and the latter is the visual appearance of that marked surface – how it looks.

Crucially, the latter is something that is there, in front of the viewer, when she faces the marked surface whose appearance is in question. If there is something odd about saying that seeing how the marked surface looks involves seeing something other than the marked surface itself, it is because the latter doesn't involve looking at anything other than the marked surface itself. The look of the surface is not something that sits on top of or alongside it; it does not occupy another place at which to direct one's gaze. But this is because the look is not the same sort of visible object as the marked surface that has it. It does not mean that it is not a distinct thing, nor does it mean that it is not *there to be seen*, in a perfectly good sense of those words.

This analysis also allows us to understand the intimate relationship between these two objects of sight, or the way in which we experience seeing one *in virtue of seeing* the other. A marked surface makes its own appearance visible, insofar as we would not be able to see the look of the surface if the surface itself were not there. Moreover, it is clearly in virtue of *seeing* the surface that we manage see how it looks. In that way their relationship is quite different from that of an opaque object and the light that illuminates it (or an opaque object and the mirror it is reflected in), even though one may see the first in virtue of the second. Recalling the point that seeing how a thing looks requires a further exercise of visual skill, beyond what is required to see it, also gives us a grip on the way that seeing the one is the basis for seeing the other. And it also gives us an alternative gloss on the special perceptual capacity to which seeing-in is indebted: it not a capacity for seeing things that are not present, but rather the capacity to see how things look.[30]

Images, in the sense I have fleshed out, present particularly ripe opportunities to exercise this capacity. Importantly, they can provide this opportunity in this way without having been *designed* to do so. That is, images can be composed of naturally occurring or haphazardly created marks. Such images are not pictures or representations of any kind; they do not represent what they are images of. Something's being an image is simply a matter of

[30] It is important to note that this way of thinking about images does not entail that all pictorial images are images of their own appearances. It simply entails that in saying what an image is of, one is characterizing its appearance, or how it looks. It also does not entail that whatever is true about what a picture depicts, the same can be said about its look. The horse in Xu's picture is galloping, but the look of its surface does not gallop. The point is just: to say that a picture depicts a galloping horse is to characterize its look, and to say that one sees a galloping horse in the picture is to say that one sees the picture and sees a certain kind of look that it has, one which is characterized in term of the idea of a galloping horse. This leaves much to be said about the use of these words to characterize the picture's look. The next chapter will approach this question; for another discussion of our characterizations of things look, see Martin 2010.

its visual potential, not the way it is put to use or is designed to function. This kind of potential allows images to be put to certain uses or designed to serve certain functions.

Having gotten a grip on the broader phenomenon of seeing-in by subsuming pictures under the heading of images, we can turn to the question of what distinguishes the pictorial species of seeing-in from others. We can do so by considering what distinguishes pictorial representations from natural images. The function that images serve when they constitute pictorial representations is that of *display*. Rather than just having particularly visible looks, they are *designed* to make these looks visible. So the experience of seeing something as depicting what it does is one of being shown how something looks.[31] But it is not only that: this experience involves understanding that one is being shown what one sees, that the thing one is looking at is intended for the purpose of display.[32]

The examples of different species of seeing-in illustrate that images can play a role in psychological diagnosis and artistic inspiration, without counting as depictions. But these instances are also closely connected to the phenomenon of depiction. In the case of the activity recommended in the *Trattato*, the point of finding images on haphazardly marked surfaces is to identify new ways of marking surfaces intentionally to create *pictorial* images. It would not be much of a stretch to think of the activity as an imaginative exercise, in which one looks at damp-stained walls and the like *as if* they contained pictures. It would also not be much of a stretch to understand the question "What might this be?" – the very vague demand of the Rorschach test – as elliptically prompting the same exercise: "What might this be a picture of?" In that sense, these visual experiences are ones of seeing pictures in things that do not actually depict – not taking them *to be* pictures, but seeing in them the potential to be depict certain things, and to put certain ways of looking on display.

Given the relationship between these experiences of images and depiction, one might wonder whether the capacity for seeing-in is a fundamentally pictorial one. If it were, Wollheim's commitment to defining depiction in terms of the experience that manifests this capacity would be less illuminating. It would say that pictures are things that afford the experience made possible by our capacity to see things as pictures. But these observations do entail that conclusion. The fact that Rorschach tests and the imaginative inspection of haphazardly marked surfaces are related to our experience of pictures does make them cases of pictorial experience. It does not rule out that they fall under a broader category of alongside the experience of depiction. And they do, if the notion of an image can be explicated in terms of the perception of visual appearance, without any reference to depiction.

I've presented a way of thinking about seeing something as depicting what it does that diverges from Wollheim's in certain ways, but despite that it upholds many of his core commitments. It maintains the idea that the experience is fundamentally a form of visual perception, in that it involves *seeing* a depicted object, not merely knowing what a picture depicts. I have also remained faithful to Wollheim's strategy for making clear distinctions between experiences according to their phenomenology. I have done so insofar as I have searched both for a way of uniting the target experience with others, as well as for a way of pointing to what makes it distinct from the others. The particulars of my results align with

---

[31] This does not entail that all pictures are made with the intention of being shown to anyone in particular, or anyone at all. It means only that they are made with the intention that they be suited to that kind of purpose.
[32] Maynard discusses this aspect of representational images as well, noting that the phenomenon "may appear complicated or paradoxical" but is "perfectly familiar," and indeed is ubiquitous in our social lives (Maynard 2000, 31).

his, in that I have connected the "genus" to a fundamental perceptual capacity, and I have differentiated the species by reference to intentions. But the intentions I have drawn on are intentions to show things, to produce things such that they make things visible and also appear designed to do so.[33] And the perceptual capacity I have pointed to is not that of seeing things that aren't present, but the capacity to see how things look.

[33] This is importantly different from the intention to produce a certain kind of experience. Thus, this view of depiction is not a version of what Hyman calls "subjectivism" (Hyman 2000, 2003, 2006).

Chapter 4
Seeing What Things Look Like: Depiction and Visual Resemblance

Fig. 2. Hokusai, *Lake Suwa in Shinano Province (Shinshu Suwako).* Woodblock print, ca. 1830-32

In the last chapter, I investigated the nature of the kind of visual experience that pictures afford. In particular, I sought to clarify the kind of richness or complexity that distinguishes it from the experience afforded by other visible things. This chapter turns, at first, from our experience of pictures to pictures themselves, to the question of what it is to be a picture. Ultimately, however, this question is bound up with the phenomenological issues of the last chapter, and it will force us to look at them again, from a different angle.

The primary question I now want to consider is whether and how the phenomenon of *visual resemblance* is implicated in depiction. Is a picture a representation that necessarily looks like what it represents? Are there any logical entailments between statements of depiction and statements of visual resemblance?

The idea that a picture must look like what it depicts has a long history. Plato's discussion of painting in Book X of the *Republic* can be read as affirming it. Plato says that the distinctive ability of a painter is to create the appearances of things – things of any kind, as long as they are visible. The pictures a painter produces resemble things of the kinds that they have chosen to depict, but because they resemble them in visual appearance alone, they do not count as instances of those kinds – they merely look like them. For Plato, the superficiality of this resemblance is the important point. He uses this conception of depiction to argue that pictures are "far removed from the truth," and that the person who makes them does not need to have any knowledge about the world in order to do so (Plato,

598b). For Plato, knowledge is of eternal Forms, and visual appearance is not a domain of knowledge. Because they are manufactured appearances, pictures appeal to the part of the soul (and the kind of person) that is unmoved by the truth, and picture-making is therefore a dangerous element of society.

If we find Plato's conclusions about the psychological and societal effects of pictures to be unacceptable, we can reject his epistemology and metaphysics, and hold on to the claim that pictures resemble the things they depict. In comparison to a division between parts of the human soul and the posit of eternal Forms as the ultimate objects of knowledge, the claim that a picture of a tree has the appearance of a tree can sound like common sense. But in fact it too is a philosophical thesis. It does not express one of the plain facts or everyday truths about pictures, but rather a particular way of understanding those facts.

My examination of this way of thinking about depiction is divided into four sections. The first will present a way of motivating the idea that a picture must resemble what it depicts, and some preliminary reasons for thinking that it is inadequate as an understanding of depiction. The motivation has to do with the contrast between depiction and linguistic expression, and the core of the objection to it is that it comes with untenable presuppositions about visual experience. The second section will address an ambiguity in our characterizations of things as pictures, which complicates the question of whether depiction requires resemblance. The third will examine attempts to defend the Platonic conception of depiction by making its claim more precise in certain ways. It will also argue that these attempts are not successful, because they too are based on unfounded claims about visual appearance. The fourth will tie these issues to considerations about the ability to see a picture as depicting what it does, and argue that a conception of depiction in terms of resemblance threatens to misconstrue this as well.

1. Pictures versus words

The claim that a picture must resemble what it depicts is *a* way of explaining our everyday distinction between pictures and linguistic representations; it is *one* way of pinpointing the contrast between them. But it does not go without saying that it is the right way. As readily as we distinguish pictures from words, it is not obvious exactly what distinction we are making. A picture is a graphic visual representation, in the sense that it is an arrangement of visible marks on a surface. But linguistic representations can also be graphic and visual in this sense. A picture's significance is dependent on the appearance of the marks that compose it. But the meaning of a written text also depends on its appearance in a certain way. It is because the marks are visibly shaped in certain ways, to form letters that are arranged to spell words, that the text on the page contains a description of one thing rather than another. Intuitively, we think that the way the meaning of the text depends on its appearance is different from the way the significance of a picture depends on its appearance. But that is not the same as having a way to articulate the difference.

But the fact that depiction and written language are both graphic and visual in some way does not mean that they are graphic and visual in the same way. If we get a bit more precise about the way in which written text is graphic and visual, the difference might become clearer. It is true that *which* words appear on a page of text depends on how the page appears to be marked. But the *meaning* of these words does not depend on the appearance of the marks that compose them. Instead, it depends on the conventions in place for using marks with that appearance to express thoughts. This might point to the way in which pictures are distinctive. Perhaps the appearance of a picture *immediately* determines its particular significance, rather than by way of conventions involving things of that

appearance. Arthur Danto expresses this conviction when he says that there is "an immediate perceptual connection between pictures and things [which they represent], as there is only in rare instances between words and things [which they represent]" (Danto 2001, 2).

But if this is correct, the distinction between depiction and linguistic expression is still not entirely clear. What is the nature of this immediate perceptual connection between a picture's appearance and its significance? The claim that a picture must resemble what it depicts provides an answer to this question. Danto arrives at this claim by considering how it is that we can come to know what a picture depicts. He considers it to be a "dividend" of our more general ability to recognize things as being of certain types by sight. Recognizing a picture of a thing of some sort, a picture of a horse, for example, involves the same capacity that recognizing a horse by sight does. That is, horses and pictures of horses alike engage our perceptual apparatus in a particular way. Danto understands this to mean that they exhibit the same kind of visual appearance. In other words, a picture of a horse must look like a horse (Danto 2001, 2-3).

As entrenched as the appeal of this philosophical view of depiction might be, it has had detractors. Wollheim is one: his analysis of depiction in terms of seeing-in is put forth as an alternative to the view that depiction depends on resemblance (Wollheim 2003). Despite that, some of Wollheim's readers have thought that his insights are compatible with the claim that pictures resemble what they depict, and have proposed that the phenomenon of seeing-in should be understood as the phenomenon of seeing resemblance (Hopkins 2003, Peacocke 1987). Instead of wading into that debate, I'll focus on Nelson Goodman, whose discussion of depiction in *Languages of Art* sets a relatively clear agenda for philosophers interested in defending an analysis of depiction in terms of resemblance. Goodman does not deny that depiction and linguistic expression are different forms of representation, but he argues that the "naïve" idea that a picture must resemble what it depicts is more misleading than promising, when it comes to understanding how they compare – indeed, he claims that "more error could hardly be compressed into so short a formula" (Goodman 1976, 4).

The main thrust of Goodman's opposition to the "naïve" idea is that it is "stopped at the start by inability to specify what is to be copied" – that is, to specify *how* a picture must resemble what it depicts (Goodman 1976, 9). All things have various aspects and facets, and no one of them has the special status of being *the* way the thing is. Even if we take the resemblance at issue to be limited to that of visual appearance, the same point applies: things' visual appearances are themselves multi-faceted, and again, no one of these facets counts as *the* way the thing looks. But resemblance in *all* respects is clearly not a requirement for depiction. So the claim must be that there is some particular respect in which a picture must visually resemble what it depicts. This requires privileging some particular respect of visual resemblance. Goodman suggests that the only justifiable way to do so would be to consider a certain circumstance and a certain frame of mind to be the correct one, for the purposes of depiction: "the normal eye, at proper range, from a favorable angle, in good light, without instrumentation, unprejudiced by affections or animosities or interests, unembellished by thought or interpretation." But, Goodman thinks, this step finally reveals that the whole idea is doomed. There is no such "innocent eye": no such thing as perception that is not "regulated by need and prejudice." Since there are no grounds for privileging any particular perspective, there can be no privileged aspect of visual appearance to appeal to, and so no way of saying how a picture must resemble what it depicts (Goodman 1976, 6-9).

This argument is quick, but it makes a powerful point: that if depiction is distinctive in virtue of the way it requires resemblance, it should be possible to specify the way in which

resemblance is at issue – what kind of resemblance is required. In order for the claim to constitute a meaningful characterization of depiction, it must refer to a specific respect in which pictures must resemble what they depict.

Hyman argues that the possibility of specifying the way in which depiction requires resemblance does not fall with the rejection of the "myth of the innocent eye" as Goodman claims, and that we can articulate a general principle that purports to pick out the kind of resemblance that depiction entails. My aim in the remainder of this chapter will be to investigate whether that is so. I will argue that the claim that depiction requires a specific kind of resemblance inevitably relies on a problematic assumption about perception. The myth of the innocent eye, and its idea of a patchwork of shapes and colors, may not be the best characterization of the problem. I will locate it instead in the assumption that previous chapters have questioned, that seeing something entails seeing how it looks.

Before presenting Hyman's principle and his argument for it, I will need to take a detour through Goodman's and Hyman's discussions of the grammar of statements of depiction. Both philosophers note that these statements are often ambiguous, and provide different ways of clarifying their various meanings. This disambiguation is necessary in order to address whether depiction entails resemblance, in that we must understand the sense of "depiction" that the claim employs in order to evaluate its truth. But these claims about the way we use "depicts" must be evaluated as well. Interestingly, Goodman and Hyman both rely on analogies between depiction and verbal expression in order to explain the kinds of significance that our statements of depiction can have. Goodman and Hyman's claims about resemblance must be understood in light of this framework, so it is important to grasp it before trying to evaluate their claims. But, as I'll argue, the use of these analogies is itself problematic. So the grammatical detour will not just be stage setting. It will also accomplish part of my criticism of Hyman and Goodman, and it will present a different suggestion about the contrast between pictures and words.

2. Statements of depiction

Both Goodman and Hyman note that "depicts" (like "is a picture of") sometimes expresses a two-place relation between the terms in subject and object position, and sometimes expresses a monadic predicate that characterizes the subject term.[34] That is, sometimes things are characterized as pictures of a certain kind in virtue of a relation they bear to something else, and sometimes they are characterized as pictures of a certain kind simply in virtue of what they themselves are or are like.

Though they agree about this fact, they disagree to some extent about how it should be illustrated. When it comes to the former type of characterization, there is less of a dispute between them. Both take the paradigm cases to be characterizations of things as portraits of individuals: statements like "This depicts the Duke of Wellington" or "This is a painting of Henry VIII" are, for both philosophers, statements that characterize the grammatical subject as a picture by relating it to a specific person.[35] Statements of depiction that contain names

[34] Other philosophers discuss this as well. Budd 2009 (16-17) gives an account very similar to Hyman's. Forbes 2006 (130-150) provides an account of the behavior of verbs of depiction as part of a formal semantic account of intentional verbs.

[35] At one point, Goodman also seems to refer to the task of drawing from life, making a picture of an object in front of you (Goodman 1976, 6). "This is a picture of my teapot" can also be understood as a statement of a relation between a picture and another thing, the thing on which the picture was based. I will discuss drawing from life in the next chapter, and argue that the relation between a drawing from life and the object on which it

of specific buildings and places are also treated this way: "This is a picture of Marlborough Castle" and "This depicts Monet's garden in Giverny," according to Hyman and Goodman, can be read as stating that the same relation holds between these subjects, and the castle and the garden, respectively. Goodman calls this relation "denotation," and considers it to be the same as the one that holds between the words "Marlborough Castle" and the castle itself. Hyman calls the relation "reference," and also considers it to be analogous to naming.

Hyman and Goodman arrive at the notion of a non-relational statement of depiction quite differently, however. Goodman gets there by noting that pictures can illustrate fictions. For example, the covers of many editions of Dickens's 1836 novel *The Pickwick Papers* contain a picture of its protagonist Samuel Pickwick. These pictures are obviously different in kind from pictures of other fictional characters, for instance the many pictures of Gusty, a feisty unicorn with a green and red mane, that appear in the 1986 animated television show *My Little Pony*. But according to Goodman, the salient difference between these two sets of pictures cannot be accounted for in terms of their denotation, since all of them have "null denotation" (Goodman 1976, 21). For him, this shows that there is a different way of classifying things as pictures of various types, other than by stating what they denote.

Of course, we can classify things that *are* pictures in all sorts of ways. We can distinguish paintings from mosaics, oil paintings from watercolors, authentic Rembrandts from forgeries, the ones that hang in the Louvre from the ones that don't. But the distinction between pictures of Pickwick and pictures of Gusty is an especially interesting kind. It has to do with what they *represent* – with the way in which they are *pictures*, as opposed to any other sort of object, which may be made by a particular person or composed of various materials. But according to Goodman, it does not have to do with what they *denote*. Goodman says this marks a "simple fact" that things can be classified as types of pictures according to what they represent, in addition to what they denote. Statements of the former sort of classification can be used to characterize *any* picture, not just pictures of fictions. They are expressed using the same locutions as those of the latter ("depicts," "pictorially represents," "is a picture of"), but it is possible to distinguish between these meanings and specify which one is intended. Statements of what something pictorially represents, rather than what it denotes, employ monadic predicates, and the terms that follow the verb serve to modify the verb – to specify a way of depicting – rather than to refer to a relatum. Goodman suggests that we make this explicit with hyphenation, calling something a Pickwick-representing-picture or a man-representing-picture instead of saying that it depicts Pickwick or is a picture of a man (Goodman 1976, 21-22). This device allows us to state more perspicuously the simple fact that there is a difference between being a man-representing-picture and being a picture that denotes a man.

Hyman thinks it is not so clear that the difference between a picture of Pickwick and a picture of Gusty the unicorn can't be accounted for by a difference in what they denote, or refer to. "Arguably," he explains, "what the referring use of a name requires is that the speaker be able to identify whom or what she is referring to." This requirement is met "quite easily" for fictional characters, since we can describe the roles they play in the works of fiction that they are part of (Hyman 2012, 134). So, arguably, the names "Pickwick" and "Gusty" can be used to refer to different things, and a picture can be considered a portrait of one rather than the other.

---

is based is distinct from that of portraiture. But it does provide, nonetheless, a way of characterizing something as a picture by relating it to something.

Hyman's route to non-relational statements of depiction focuses instead on statements that contain indefinite noun phrases: "This is a picture of a horse," or "This is a picture of a unicorn." He notes that a picture may be of *a* horse, though there is no particular horse (neither historical nor fictional) to point to as the one to which it refers. That is, something can be a picture of a horse without being a portrait of a horse. Hyman says that a picture of this sort is a "genre picture," or a picture with "generic content," as opposed to a portrait (Hyman 2012, 135). There is nothing to which it refers, but nonetheless it is a picture, and it is possible to characterize it by saying what it represents, using an indefinite noun phrase.

Goodman, it seems, would dispute the claim that pictures that are not portraits of specific things therefore do not relate to anything in the way a paradigmatic portrait relates to the person it portrays (what Hyman calls "reference" and Goodman calls "denotation"). According to him a picture, "like a predicate, may denote severally the members of a given class." He claims that a picture that illustrates a dictionary definition is "often such a representation, not denoting uniquely some one eagle, or collectively the class of eagles, but distributively eagles in general" (Goodman 1976, 21). So the fact that a picture of an eagle is not a portrait of some specific eagle (that there is no "namely rider" that can augment this statement, salva veritate) does not entail that it has no reference or denotation.

Despite these disagreements, Hyman agrees with Goodman about the "simple fact" that things can be classified as pictures of various types according to what they are related to as representations, but also independently of such consideration. The important contrast, for Hyman, is not between portraits and genre pictures, but rather between two different sorts of things that could be said of a *single* picture. Whether or not a picture is a portrait or refers to anything, we can attribute some "generic content" to it. When we do, we are making a non-relational statement about the thing in question, using a one-place predicate to characterize it. When we point to a portrait of Seabiscuit and say that it is a picture of a horse, we may not be saying that there is a thing to which it is related in a certain way, and characterizing that thing as a horse. We may be characterizing the picture itself, without considering any relations it might stand in.

Hyman goes somewhat further than Goodman does in explaining what way of characterizing pictures this is. He does so by invoking the Fregean distinction between sense and reference. These notions apply originally to linguistic expression, to "distinguish between the object that an expression stands for or designates, and the way in which the expression presents that object, the 'mode of presentation'" (Hyman 2012, 136).[36] This formulation might seem to imply that in order for an expression to have a sense, it must refer an object. It might seem to imply this because it states that an expression's sense is the way in which it presents "that object" – the object to which it refers. But this does not as Hyman understands it, there is no requirement that an expression refer to anything in order to have sense. He gives the phrase "the present king of France" as an example of an expression that has a sense, but does not refer to anything. As Hyman wants us to think of it, an expression's sense is a *way* of referring, in general. An expression might refer in a particular way, or exhibit a mode of presentation, without actually referring to anything in

[36] Goodman does note that there is a "close parallel" between the difference between being a man-representing-picture and being a picture denoting a man, and the difference between "a man-description (or man-term) and a description of (or for) a man" (Goodman 1976, 23). The way in which I consider Hyman's explanation to go further than Goodman's is that he attempts to say something about what this parallel consists in, by using the concepts of sense and reference.

that way. And its sense can be stated without reference to *which* thing or things it refers to, if any.

Hyman thinks that there is an analogous distinction to be made for pictures, and that it can help us to understand the difference between relational and non-relational statements of depiction. He illustrates the point by comparing two different portraits of Tolstoy: an 1873 Kramskoy that presents him as "dark-haired and seated," and a 1901 Repin that presents him as "grey-bearded and standing" (Hyman 2012 137). The two pictures have the same reference – Tolstoy. But the fact that one is a picture of a seated man with dark hair, while the other is a picture of a grey-bearded man standing, marks a difference with respect to their senses. We use the word "depicts" to talk about both of these aspects of depiction, and so our talk about pictures can be ambiguous. Hyman suggests we clarify this ambiguity using subscripts: "depicts$_r$" can be used to express the relation, and "depicts$_s$" can be used to express the non-relational phenomenon.

With the subtleties of these two ways of understanding our characterizations of things as pictures on the table, I want to question the parallel between depiction and linguistic expression that they both depend on. In both instances, it begins with an assimilation of the phenomenon of portraiture to the phenomenon of denotation, or the relation of reference. But the conception of the latter is part of a larger framework that is designed to organize our thought about language, and portraiture does not fill its role comfortably. My argument for this is as follows. Denotation and reference, as Goodman and Frege conceive of them, admit of two different types. Names and predicates denote in different ways; reference is either to objects or concepts. But there is no obvious analogue for the distinction between names and predicates, or objects and concepts, when it comes to depiction.

Goodman claims that there is when he says that a picture, "like a predicate, may denote severally the members of a given class" (Goodman 1976, 21). But his alleged illustration of this, a picture of an eagle accompanying the definition of "eagle" in a dictionary, is not convincing. We can accept that such a picture does stand in a relation to the predicate "is an eagle," and that it does so in virtue of what it depicts. But the fact that it appears next to the definition and supplements it does not mean that it stands for what the predicate "(is an) eagle" stands for, and denotes "severally" in the way that a predicate does. We do not describe it as a picture of eagles in general, or as a picture "for *eagle*"; it is a picture that represents *a* thing *as* an eagle. A picture of an eagle can indeed supplement a verbal definition of the word "eagle" when placed alongside it. But that is a purpose to which it can be put, *as a picture of an eagle*. If the picture were equivalent to the predicate, then we should expect the verbal definition to define it as well, instead of being supplemented by it.

By explaining statements of depiction of the form "This depicts an X" as characterizations of a picture's sense, rather than statements of its reference, Hyman implies that pictorial reference is always to an individual. But his way of talking about pictorial sense sometimes suggests that it is a pictorial analogue for predicates. For instance, he criticizes Hopkins for using the example of a picture of "a horse, but no horse in particular" to show that a picture can "depict what does not exist." Hyman claims that "on the contrary," such a picture "depicts something, a kind of animal, that *does* exist, unlike a picture of a centaur, for example" (Hyman 2012, 135). The claim that a picture of a horse depicts *a kind of animal* might seem to suggest that it is akin to a predicate, which applies to animals of a certain kind. Saying that a picture has "generic content" also encourages the idea that its significance is general in the way a predicate is. But if Xu's *Running Horse* does not portray a particular

horse, that does not make it a picture of being a horse, or of horses in general. It means that it depicts something (though nothing in particular) as being a horse.

I now want to consider the problem from a different angle. The issue so far has been about whether there are any pictures that operate like predicates, by corresponding in an analogous way to concepts. But concepts and objects, for Frege, are the referents of the constituents of sentences. This makes the attempt to construe some pictures as like predicates and others as like singular terms seem misguided: if we are going to analogize depiction and linguistic expression using these notions, the distinction should apply to the constituents of pictures, rather than sorting pictures themselves. But there is no clear way of identifying constituents of pictures as general rather than singular, either.

Frege's distinction between concepts and objects cuts across his distinction between sense and reference. But the two distinctions are importantly linked. Linguistic expressions have senses that compose full thoughts, which can be asserted and evaluated for truth or falsity. But here too there is a problem, since there is no clear systematic application for the notion of truth or assertion to depiction. Pictures can be evaluated for accuracy, but that does not mean that a picture itself can be true or false. When we ask whether Repin's portrait of Tolstoy is accurate, we are asking whether the picture depicts something as being the way Tolstoy in fact was (at a certain point) – a grey-bearded man. Answering in the affirmative means that this statement is true, not that the picture is true. It entails only that the picture is accurate, since that is what the statement says.[37] [38] [39]

For these reasons, we should be wary of the idea that portraiture constitutes a pictorial version of the reference of singular terms in a language. Hyman claims that "the most important mistake philosophers have made about depiction is to confuse or amalgamate theories about the sense of works of art and theories about their reference, or to assume that a theory of depiction is first and foremost a theory of reference – as it were, a theory of the portrait – and a theory of sense can be developed from it, rather as Wittgenstein's theory of meaning in the *Tractatus* was developed from his conception of a name" (Hyman 2012, 137-138). But I am suggesting that the mistake runs deeper; that it is a mistake to think of portraiture as a kind of reference, and it is a mistake to think that what it is to be a picture is to have the potential to refer. Portraiture is a very common purpose of depiction, but that does not make it an aspect of depiction itself. Portraits are pictures that

---

[37] It should be acknowledged that philosophers have denied this. Flint Schier advocates for the view that pictures do have truth values, despite the fact that there is no pictorial analogue of names and predicates. He puts this by stating that pictures are "saturated" symbols, borrowing another Fregean term. In the linguistic case, "saturation" is a name for the way that a "saturated" singular term combines with an "unsaturated" predicate to determine a truth value. Schier does not address the fact that names as well as statements are saturated, such that saturation does not correspond immediately to having a truth value (Schier 1986, 170). Alex Grzankowski argues that because it is possible to point to a picture and say "This is not how it was," pictures have content which can be negated, and so must have propositional content, like sentences do (Grzankowski 2015). I do not find this convincing. In that statement, "this" refers to the way things are represented by the picture, and the statement negates the claim that this was the way things actually were. It does not negate the way things are represented by the picture.

[38] I do not mean to deny that there are any legitimate or illuminating ways of applying some notion of truth to depiction or visual art in general. What I say in this chapter is compatible, for example, with Michael Podro's claim that pictorial art exhibits what he calls "truthfulness" (Podro 2010). My claim is only that it does not make sense to consider pictures to be true in the way that assertions can be.

[39] Andrew Harrison is making a related point when he says that picture-making and story-telling are fundamentally different, and that though we use pictures to illustrate stories, "*by themselves* pictures do *not* tell stories; there is nothing in the pictorial record itself that can correspond to those words of consequence, of likelihood or possibility that articulates the telling of a tale" (Harrison 2001, 43).

stand in a certain relation to particular things, in virtue of what they depict. But that does not mean that they stand *for* anything, as words do. We would do better to think of portraiture as a purpose to which pictures can be put – an alternative to that of a picture accompanying a dictionary definition. We might think of that purpose as showing what that individual looks like, or capturing its likeness.[40] To do so is not itself a matter of depicting, but of presenting a depiction in a particular way.

This means that we should be equally wary of applying the Fregean notion of sense to the phenomenon of depiction. Pictures can be characterized by saying how they represent something. But if sense is understood as a way of representing a *referent*, then the way a picture represents something is not a matter of its having a certain sense. That is not to deny that there is a distinction between relational and non-relational uses of locutions like "depicts" and "is a picture of," nor that some uses of these locutions characterize something as a portrait of an individual. Nor is it to deny that it is important to distinguish these two uses when considering the role of resemblance in depiction. But it means that we should not understand pictorial classification, or the non-relational characterizations of the way in which a picture represents something, as a matter of their having a Fregean sense.

It is important to separate these views of the distinction between relational and non-relational characterizations of things as pictures from the distinction itself. The distinction itself is important to keep in mind when turning to the question of the relationship between depiction and resemblance, which I will do in the next section. I'll begin by explaining how this distinction figures in Goodman's criticisms of the idea that resemblance can account for the nature of depiction, which were summarized in the first section. I'll then consider two responses to it: one by Budd, and one by Hyman, both of which claim that it is the non-relational phenomenon that is properly understood in terms of resemblance.

3. The role of resemblance in depiction

Goodman's treatment of the claim that a picture must resemble what it depicts initially focuses on whether what he calls pictorial denotation requires resemblance. As explained in the previous section, he thinks that denotation can be singular or general. His reasons for rejecting the claim that a picture must resemble what it denotes, which I summarized in the introduction of this chapter, are presented in a way that focuses on singular denotation. Goodman asks in what way a picture is supposed to resemble *the object* it denotes, claiming that this needs an answer if resemblance is supposed to explain what pictorial denotation involves, and argues that no answer can be given. But the issue is meant to arise for denotation generally, whether it is singular or "several." If a picture denotes objects of a certain kind, then the question is in what way it must resemble objects of that kind.

---

[40] I'll say more about portraiture in the next chapter. For a full analysis of it, see Freeland 2010. Her analysis would exclude many of the cases that Hyman considers to be portraits in an extended sense (pictures of specific places, events, or things), since it defines a portrait as a certain kind of picture of a person. It would be interesting to assess whether it provides any insight into what could be the basis for this kind of extension, but that is outside the scope of this chapter. I am not denying that there is any interesting or illuminating extension of the concept of portraiture to include the cases Hyman has in mind, but only that it is not obviously or easily assimilated to the phenomenon of reference. It would be interesting and fruitful, I think, to analyze the similarities and differences between the many uses of pictures that place them (as pictures) in relation to other things. This would involve a fuller examination of the tradition of portraiture, as well landscape painting, religious iconography, the illustration of fictional and non-fictional texts, the use of pictures in taxonomy, as well as the many uses of photographic images. Such an investigation might draw on Noë's presentation of the view that pictures of people should be thought of as "stand-ins" or "substitutes" for the people they are of (Noë 2012, 2015). I regret that there isn't room for this investigation here.

Hyman and Budd concede the point that portraiture and analogous relations do not require resemblance and are not to be analyzed in terms of it. Hyman claims that the relation between a portrait and its subject is determined by the intentions of the artist, not any kind or degree of resemblance between the two (Hyman 2012, 140). Budd also claims that it is dependent on the artist's intentions, or the causal role that a thing plays in the production of a picture. But Hyman and Budd both think that non-relational statements of depiction are about a phenomenon that does crucially involve resemblance. Budd claims, moreover, that this is the kind of statement that a theory of depiction should focus on. He argues that "an account of what is distinctive of depiction does not need to characterize the relation in which a relational picture stands to its subject," because that relation is not distinctive of depiction: "pictorial reference is not different in kind from other forms of reference" (Budd 2009, 217). A theory of depiction should not tell us what reference is, but rather what it is for something to refer *pictorially*. And this is captured not by statements of pictorial relation, but by statements of pictorial classification.

If my criticism of the concept of pictorial reference is correct, then Budd's claim can't be exactly right as stated. But even if my criticism is accepted, a point quite close to Budd's could nonetheless be made. It would be that relational statements of depiction are about relations things can stand in *given* that they are pictures of certain kinds, whereas non-relational statements are simply about what they are, as pictures. This supports the point that a philosophical analysis of depiction itself should view the non-relational phenomenon as fundamental.

Goodman does not explicitly consider whether resemblance of any kind is entailed by what he calls pictorial classification, or non-relational characterizations of depiction. He notes that the possibility of pictures with "null denotation" – pictures of fictions – provides a "further beating" for the claim that resemblance is required for depiction, in that "where a representation does not represent anything there can be no question of resemblance to what it represents" (Goodman 1976, 25). What he means is that there can be no question of resemblance to what the picture *denotes*, if the picture has no denotation. But a picture of a fictional entity is still a picture of a fictional entity of one kind rather than another. It may be a unicorn-representing-picture, or a man-representing-picture. What Goodman says seems to leave unaddressed the question of whether its counting as one rather than the other means it must resemble anything in any particular way.

Hyman claims that in order to grasp how the non-relational phenomenon of depiction involves resemblance, it is important to recognize an ambiguity in statements of resemblance. He explains the ambiguity in "resembles" by comparing it to the verb "to be." The word "is" sometimes expresses an identity, and in other cases functions as a copula, serving to characterize its subject in conjunction with a predicate. Analogously, some statements of resemblance express a relation of similarity, while in other cases, "resembles" or "is like" functions as a copula, serving to characterize a single thing. The statement "Soho is like a village" is an example of the use of "is like" as a copula. It does not mention any villages, or any particular other than Soho. It simply characterizes Soho, "by saying what it is like, rather than what it is" (Hyman 2012, 134). Hyman claims that statements of *visual* resemblance, which describe what things *look* like, can function this way as well. When we say that the cranes at the Port of Oakland look like unicorns, we are not comparing the cranes to any particular individual or individuals (after all, there are no unicorns to pick out for that purpose, or any other). Nonetheless we are making a perfectly intelligible and informative claim about these things, in saying what these things look like, or what they resemble in visual appearance.

I explained in the first section that Goodman's objections to the idea that a picture must resemble what it depicts set the agenda for defending that claim, by raising the issue of pinpointing *how* a picture necessarily resembles what it depicts. These two points, that the non-relational phenomenon of depiction is at issue, and that there is a non-relational sense of "resembles," set the stage for the kind of reply that Hyman and Budd want to make. They make clear that the claim is not about a comparison between two specific things, a picture and something else. Rather, it is about how a single thing can be characterized. Having clarified this much, they then attempt to do the specification called for, and articulate the particular way in which every picture resembles what it depicts.

Budd and Hyman pick out essentially the same feature of things, but they end up making subtly different claims about how this feature is involved in depiction.[41] I'll present each of them in turn, and assess them separately.

Budd begins with the thought that a picture is "essentially a two-dimensional representation of a three-dimensional world, in that the picture's depicted scene is visible in its two-dimensional surface" (Budd 2009, 219). Because a picture's surface is not and does not look like it extends in three-dimensions, but what it depicts can be characterized by relations in three dimensions of space, he reasons that "the only relevant sense in which a picture, seen as a depiction of its subject, can look like its subject is with respect to the two-dimensional aspect of the subject's appearance" (Budd 2009, 219).

To make clear what the two-dimensional aspect of appearance amounts to, Budd uses the idea of a description of the "visual field" associated with the experience of seeing a particular kind of scene. This is a matter of giving a "partial account of how [one's] visual experience represents the world as being" (Budd 2009, 221), by which he means a partial account of the features of visible objects. It is what remains if we disregard "distance outward" along a line of sight, and describe the directions and distances of things in space only in terms of whether they are above or below, to the right or to the left of each other (Budd 2009, 220). For any particular visible object, its "visual field shape" will be determined by its overall shape and its position relative to a line of sight.[42]

The idea of abstracting from three dimensions to two is perfectly intelligible. We can characterize visible things in terms of how they extend in two dimensions, noting how their parts (including their edges or contours) lie above, below, and to either side of each other. But we must be careful about how this relates to visual experience. To call such a characterization of a visible object or scene a "partial account of how my visual experience represents the world as being," if I am looking at a scene that fits the relevant description, is misleading. It might suggest that the two-dimensional aspect of these things is necessarily among the things I see, if I see the objects that have this aspect. But I might see a die, and see it as having a cubical shape, without seeing how, in virtue of its cubical shape and its orientation, its vertices and edges relate to each other in two dimensions. So we must keep in mind that the two-dimensional aspect of things' features is an abstraction from the unified whole that they constitute, not a necessarily manifest element of any given experience of

---

[41] The feature that Budd and Hyman pick out is also equivalent to Hopkins's "outline shape" (Hopkins 1998). Hopkins uses it to articulate a subtly different thesis, that depiction is based on *experienced* resemblance in outline shape. I won't present his view here, but many of the points I make about Budd's and Hyman's views would be relevant to an assessment of it.

[42] Though shape in the visual field is introduced by way of a contrast with three-dimensional shape, I say "overall" shape rather than "three-dimensional" shape here, because not all visible things have three-dimensional shapes. The face of a nickel, for instance, has a two-dimensional shape – it is circular. This two-dimensional shape is still distinct from its shape in the visual field relative to a given line of sight.

seeing them.

Budd uses this understanding of the two-dimensional aspect of things' visual appearances to articulate a thesis about how pictures must resemble what they depict: they must resemble them with respect to the two-dimensional aspect of their appearance (Budd 2009, 222). Budd later articulates the point in terms of isomorphism: a picture is a two-dimensional expanse whose markings give it a certain structure, and in virtue of this it is isomorphic to the scene or object it depicts. This isomorphism is a visual one, in that this structure is part of the visual appearance of the surface, and it can strike viewers as similar to other sorts of things with that structure.

Budd's clarification of the concept of the two-dimensional aspect of things' visual appearances doesn't prove the claim in question on its own. That only shows that pictures *can* resemble what they depict with respect to two-dimensional aspect of things' visual appearances. The fact that pictures appear on the surfaces of things may mean that *if* they resemble what they depict, they must do so with respect to the two-dimensional aspect of things' visual appearances – at least, they must do so with respect to features that two-dimensional things can have. But the reason for accepting the antecedent of that conditional must be found elsewhere.

We might think that the reason is simply the prima facie plausibility that the claim has. It expresses a longstanding and widespread conception of depiction, and perhaps, in the absence of reasons to deny it, we should grant that it is true. Budd does not seem to think that we can leave the matter there, however. He provides abductive support for his thesis, claiming that it "accounts for many of the most significant and distinctive features of pictorial representation" (Budd 2009, 235). I will return to the considerations he presents in the last section. For now I want to note that this way of motivating the thesis ignores the question of whether there are any reasons to deny it. It might seem that doing so is fair, at this point in the investigation, since Goodman's criticisms have been addressed. But I think there is a reason to deny it that is in line with the spirit of Goodman's criticisms, even if it is not explicitly covered by his discussion.

The problem, which becomes more vivid once the thesis is made more precise by reference to the two-dimensional aspect of appearance, is that it requires that there be a characteristic two-dimensional aspect of anything that can be depicted. It states that what it is for something to be a picture of an X, for any X, is for it to resemble an X with respect to the two-dimensional aspect of visual appearance. But in order for something to resemble an X in this way, it must be possible to specify how Xs look, with respect to this feature. Not how an X could look, or would look in certain circumstances, in this respect – the claim is that the picture must look the way Xs do, in general. But there are many values of X, that is, ways for something to be depicted, for which there is no particular corresponding structure of the visual field.

Take, for example, pictures of horses. It is true that horses are not perfectly spherical, and perhaps it is also true that no fictional horse could be perfectly spherical either. But it does not follow from the fact that it is impossible for a horse to have any number of shapes that there is a particular type of two-dimensional shape that is characteristic of horses' visual appearances, such that anything with that shape will resemble a horse in that respect.

Hyman's claim about depiction and resemblance carefully circumvents this issue. But, I think, it is not on better footing as a result. I'll turn to his view now.

Hyman draws our attention to the two-dimensional aspect of visual appearance in terms of a feature he calls "occlusion shape." He explains that this is the property a visible

thing has in virtue of being hypothetically occluded (that is, of being occlud*able*) by an "opaque patch" of a particular shape (Hyman 2006, 76). Occlusion shape is relative to a "line of sight," in that what is relevant to an object's occlusion shape is occlusion *from view* (rather than occlusion in the sense of blocking a passageway).[43] But that does not mean that it is subjective, in the sense of being a feature of experience rather than a feature of perceptible objects in the world (Hyman 2012, 143). In fact, it is a feature of *all* visible objects. In this sense, it is a fundamental aspect of the structure of the visible world. (Despite that, it is not a feature that things are necessarily seen as having.) Though occlusion shape is specified in two dimensions, it can nonetheless be attributed to objects that are extended in three dimensions of space. Moreover, though two-dimensional visible things have occlusion shapes, they are not necessarily congruent. The square face of a six-sided die, for instance, would be occluded by a parallelogram-shaped patch relative to various lines of sight. In sum, we can think of occlusion shapes as defining aspects of the way things' shapes appear visually.

According to Hyman, the concept of occlusion shape allows us to see that there is a systematic relationship between the shapes demarcated on a picture's surface and the way it depicts something. He articulates that relationship in the following way: "the shape of the region on a picture's surface is the same as the occlusion shape of the object it represents." This proves, he thinks, that there is a specific way in which a picture must resemble (non-relationally) what it depicts: it means that there is "an exact resemblance" between the shapes on a picture's surface and the occlusion shapes those regions depict (Hyman 2012, 143).[44]

This thesis differs from Budd's in that it does not say that *whatever* a picture depicts, it resembles it in the specified respect. Instead, it requires that the description of what is depicted be given in terms of occlusion shape. Whatever occlusion shape is depicted by a part of a picture, the picture delineates a region of that two-dimensional shape. This might not be obvious from Hyman's statement of the "shape-rule for pictures," but it can be seen if we consider what he presents as an easy way of proving that the rule applies. The proof involves a thought experiment:

> The experiment is to try to trace the shape of the part of a picture which depicts something – say, a house or a tree or a man or a part of his body – by running a finger across its surface, without simultaneously tracing the occlusion shape of the corresponding part of the picture's internal subject – the house or the tree or the man in the picture. (Hyman 2000, 26)

Hyman thinks that it is clear that this cannot be done: in tracing the part of the picture that depicts something, one will necessarily be tracing *the occlusion shape of the corresponding part of the picture's internal subject* – that is, the occlusion shape that this part of the picture represents something as having. Thus, any picture must depict something as having an occlusion shape. And it is the depiction of occlusion shape, and nothing else, that entails resemblance of a particular sort. Hyman considers this shape-rule to be the "defensible residue" of the

---

[43] Occluding from view is not equivalent to blocking the path of light. As Kalderon points out, highlights on shiny surfaces and halos around luminous objects can occlude things from view (Kalderon 2015, 81), but they don't block the path of light.

[44] Hyman considers this to be a relationship between the shape of a part of a picture and its sense, but it is possible to state the point without using the notion of sense. It is important, of course, that the statement employs the non-relational use of "depicts."

conception of depiction we get from Plato; he thinks it is the truth that remains once we recognize that only the "basic representation of visible objects" can be explained in terms of a requirement of resemblance between picture and depicted object (Hyman 2006, 71).

But it is not clear that the shape-rule is so easily proven. The description of this experiment refers to "tracing the occlusion shape of the house or the man or the tree in the picture," and suggests we should be able to compare the result of this task with a tracing of the part of the picture that depicts the house or man or tree. But it is not obvious what the former task requires. Something can be traced with a finger insofar as it has a defined boundary that the movement of a finger can follow. A part of a picture can be traced in this way, if it is a region of a surface that is visibly defined or demarcated in some way. Tracing it like this can make the shape of this region more perspicuous than it might be otherwise, to oneself and to any onlookers. But a depicted occlusion shape cannot be traced in the same way. It does not have its own independent boundary for us to follow with a finger (remember that it is not to be thought of as something independent of the picture, to which it is related). Tracing a part of a picture can serve the purpose of making perspicuous what it depicts, however. To show that there are two rowers in the boat depicted in Hokusai's print reproduced on the first page of this chapter, I might trace the part of the page that depicts them, and thus make these figures more perspicuous to you. This may be what Hyman has in mind, and why he thinks that in tracing part of a picture one can simultaneously trace the occlusion shape of a depicted object.

But one can make an occlusion shape perspicuous in this way only if the picture depicts something as having an occlusion shape in the first place. Hyman supports the claim that occlusion shape is an aspect of things that pictures "invariably include" by observing that "we cannot discover different aspects of an object represented in a picture by moving around it and studying it from different angles" (Hyman 2012, 142). This is supposed to entail that what a picture depicts, when "expressed in the most general terms," is "an aspect or view of an object or arrangement of objects – or several aspects or views, in unusual cases… – relative to a line (or lines) of sight" (Hyman 2012, 142). But the observation does not secure this conclusion. We can explain the observation simply by noting that pictures are marked surfaces, without making any claims about what they can *depict*. Given that a picture is a surface, the angle from which we can study it (for whatever purpose) is limited. We can't study a surface from behind, or from an extremely oblique angle, since the object whose surface it is will block it from our view.

Surfaces do not have, but rather *are* two-dimensional aspects of things. Demarcated surfaces can make two-dimensional shapes particularly salient to us, in virtue of the visible discontinuities of which they are composed. These discontinuities can be seen as defining the boundaries of two-dimensional expanses of various shapes. Given that pictures are visibly marked surfaces, their visual appearance can be characterized using two-dimensional geometrical concepts. But this alone does not entail that what they depict can necessarily be characterized in that way.

It might seem that the depiction of two-dimensional shape is secured by the supposed fact that pictures necessarily depict visible things. Visible things necessarily have two-dimensional aspects, and so what a picture depicts must have some two-dimensional aspect. But the fact that any particular thing must have some two-dimensional shape doesn't mean that it must be depicted as having a two-dimensional shape. Nor does it mean that a picture of a thing must depict it as having a two-dimensional shape.

We must conclude that the facts of optics do not justify the claim that any picture depicts occlusion shape. Without any other reason to accept that depiction is necessarily the

depiction of occlusion shape, Hyman's thought experiment cannot prove his point. Tracing a part of a picture that depicts a given thing cannot be relied on to identify a depicted occlusion shape. It can only be relied on to identify the shape of that part of the picture. The fact that we can trace the part of a picture that depicts something does seem to indicate something about depiction: that it depends on the way that visibly demarcated regions of a surface are shaped. This is, arguably, something interesting and distinctive about depiction. Though we can locate the part of a text that describes something, and we can analyze the text into sections that are about a series of different things, in doing so we are not demarcating regions of anything in a way that inherently makes their shapes visually salient. This vindicates the claim that pictures depict in virtue of their visible form.[45] But it does not mean that there is any "exact resemblance" between the shapes that compose a picture and anything it depicts, or any "strict and invariable relationship between the shapes and colors on a picture's surface and the objects it depicts" (Hyman 2006, 73).

In attempting to question Budd's and Hyman's views, I hope to have made clear some important entailments of these ways of thinking about depiction. Budd's requires that there be some characteristic way of appearing for anything that can be depicted, while Hyman's requires that there is a certain kind of description that can be given of what any picture depicts. I have tried to show why the particular requirements of these two views are not ones we should accept. I think that both of these kinds of requirements are problematic in general, and I hope that these criticisms could contribute to an understanding of why that is, though I have not spelled it out here. In the next and last section, I will attempt to make a little more progress towards this understanding from a different angle, by connecting these ways of thinking about depiction to ideas about what is involved in the experience of seeing things as pictures of various kinds.

4. Showing and reminding

In the previous section, I mentioned that Budd defends his thesis by an abductive argument, that certain central features of depiction are best explained by his account. I want to focus on the second of the ten considerations he lists. It is the claim is that "a spectator who has no idea what a certain state of affairs would look like (or – to cover fictional kinds – is supposed to look [like]) cannot see a picture as a depiction of that state of affairs" (Budd 2009, 235). This claim does indeed seem to dovetail neatly with the thesis that a picture must resemble what it depicts with respect to a certain aspect of visual appearance. If part of what it is to depict an X is to look like it in a specific way, then to see something as depicting what it does, one must see it as looking like an X in that way. Having no idea what an X looks like would rule that out.

But if we are to take this to be an uncontroversial fact about depiction, and an explanandum or starting point for any analysis of the phenomenon, we must be careful about how this claim is interpreted. It is natural to think of seeing something as looking like an X as a matter of its being reminiscent of an X – such that the experience is one of being reminded of a certain kind of visual appearance with which one is familiar, and associates

[45] Though I am granting the point here, I am not sure that this kind of tracing actually can be accomplished for every depiction of an object. Some depictions of objects are constructed out of marks whose salient features are one-dimensional, rather than two-dimensional. They can be traced, in the sense that these marks can be followed and reproduced. But the result in their case is not an outline of a region of a two-dimensional surface (the shape that could be had by a "patch" that could occlude something); instead, it is a set of lines crossing a two-dimensional surface. For an interesting discussion of the uses of one- and two-dimensional features of marks in depiction, see Maynard 2005.

with Xs. This may encourage the thought that for any kind of thing we can see, there is a particular kind of appearance that we associate with it, and pave the way for the idea that depiction involves exploiting these associations. In "On Looking At a Picture," Budd seems to take this tack. He starts with the intuition that "whatever a picture depicts, you would not see it as a depiction of that thing if you were unaware of what that thing looks like" (Budd 2008, 204). From there, he takes a "short step to the conclusion that it is in virtue of your knowledge of how something looks that you are able to see a picture as a depiction of that thing" (Budd 2008, 204).

But the short step is an abductive one, and the conclusion is not the only explanation available. It would also be explained by the supposition that pictures are designed to show us what things look like, and that in looking at a picture and seeing it as depicting what it does, one is *made* aware of the look of what it depicts. That too would entail that "whatever a picture depicts, you would not see it as a depiction of that thing if you were unaware of what that thing looks like" (Budd 2008, 204).

Budd may overlook this possibility because he is conflating certain candidate preconditions for seeing something as depicting a thing of a certain kind. He may have in mind that one can't see a picture as depicting an X unless one knows *what an X is*, and unless one can *recognize an X by sight* (unless one "knows an x when one sees one"). But knowing what an x is and being able to recognize one by sight does not amount to knowing how an X looks. (I have argued for this in Chapter 2.) So both knowing what an X is and being able to recognize one by sight may be required for seeing something as depicting an X, but that does not entail that it also requires knowing what an X looks like.

Budd does leave open the possibility of learning from a picture what something looks like, in a certain way. For him this would have to involve seeing the picture as exhibiting a certain look, and being independently informed that it is a picture of an X, and then putting together on that basis that the look of the picture is that of an X. But this is possible only given that one sees the picture in a suitable way at the outset – as depicting something with the look that an X in fact has. What puts a viewer in a position to do that? Is there some antecedent familiarity with the looks of things that the viewer must have in order to see the picture as exhibiting this specific way of looking? It is not clear what the content of that knowledge would be, nor is it clear that the viewer needs any such knowledge. We can explain how the viewer sees the relevant look exhibited by the picture by noting the simple fact that she has a clear view of the thing that exhibits it – the picture itself. Of course, the viewer may need to have and draw on a variety of visual skills and abilities to do this, but they need not amount to knowledge of or familiarity with any particular kind of visual appearance.

This means that at some level, pictures must be able to show us things we haven't seen before. When we look at them, we can see how the things they depict *look*, without already being familiar with these ways of looking. But if it is possible to get this far with a picture without any antecedent knowledge of how things look, why not think one can get all the way to seeing the picture as depicting what it does? This does not contradict the claim that if one *does* see the picture as depicting what it does, then one is thereby aware of how such a thing looks. It simply accounts for it in a different way, by emphasizing the role that pictures play in enabling us to see, rather than reminding us of things we've seen before. Knowledge of how things look can be thought of as an upshot, rather than a precondition, of seeing things depicted.

Above, I distinguished this from other candidate preconditions for seeing something as depicting an X, namely knowing what an X is, and being able to recognize an X by sight.

It is worth noting that these might not apply universally, and in some cases might be the upshots of seeing things depicted. As Goodman points out, one can learn about a certain type of thing through being shown pictures, without ever having seen or heard about things of that type in any other context (Goodman 1976, 24-25). A person might acquire the concept of a unicorn by being shown pictures of unicorns, and being told nothing more about each one than, "This is a picture of a unicorn" (or as Goodman would prefer to put it, "This is a unicorn-representing-picture"). Again, this does not contradict the observation that being able to see something as a picture of an X goes hand in hand with knowing what an X is, knowing an X by sight, and knowing what an X looks like. It just means that they can *come* hand in hand, as the spoils of experience with pictures.

Insofar as resemblance is tied to reminiscence, the conception of depiction that I have been questioning leads us to think of pictures as providing a kind of visual experience that their viewers have had before. It also encourages a view of visual experience that I've questioned in previous chapters, according to which whenever we see things, we see how they look. Hyman and Budd are careful to reject the "myth of the innocent eye," the idea that we directly or strictly see only a two-dimensional patchwork of colors and shapes. They both maintain that we generally see things as arrayed in three dimensions, and that the appearances of things, including their two-dimensional aspects, in no way block our view of the things that have them. But their views do not make room for the way in which pictures can show us how things look, and so they encourage us to think of the looks of things (and in particular, their two-dimensional aspects) as immediately manifest, albeit along with the things that have them.

It is important to be clear about what we do not have to deny in rejecting this conception of depiction, and how much of the spirit of Budd's and Hyman's discussions can remain in place. We do not have to deny that there is an intimate connection between visual experience and depiction, or that our ability to make and use pictures is predicated on our ability to see. Nor do we have to deny that pictures depict what they do because they contain visible demarcations of two-dimensional shapes, which are aspects of the appearances of three-dimensional things. If we hold onto all this while giving up the idea that depiction is based on resemblance and deny the assumptions that it requires, we can think differently about the role of two-dimensional shape in depiction and in vision. If we recognize occlusion as fundamental to vision and understand its conceptual tie to two-dimensional shape, but stop taking for granted that this is automatically manifest to viewers, we can understand better what it is to see the two-dimensional aspect of visual appearance, and appreciate what it takes to do so. Then we will be in a position to understand depiction as a form of representation that shows us how things look.

Chapter 5
Drawing and Discovery

The previous chapters have addressed issues surrounding the visual experience of both pictures and non-pictorial objects. In the first two, I discussed the way in which seeing requires skill, and the relationship between the skill involved in seeing a given object and that involved in seeing how it looks. The third and fourth chapters turned to pictures and the perception of them, and identified ways in which their looks are all-important to their pictorial status. In this final chapter, I'll make one more attempt to explore the relationship between pictures and vision, by thinking about how the visible objects serve as the input or basis for a certain kind of depiction. I'll end with some comments about pictorial realism.

The topic of this chapter introduces a new kind of distinction between pictures. The last chapter noted that some pictures portray specific things – people, places, animals, inanimate objects – while others do not. Even those that do not can be characterized as depicting something, or as pictures of things. But in characterizing pictures in this way, these locutions are not used to state a relation between the grammatical subject and something else; rather, they are used to form one-place predicates that characterize the subject as a particular kind of picture. I argued that the relation between a picture and whatever it portrays should not be understood on the model of linguistic reference, and that the distinction between the portrayal of specific things and the classification as a picture of a thing of a certain type should not be understood as contrasting aspects of depiction in the way that reference and sense are contrasting aspects of linguistic meaning. That is not to deny that we can characterize things as pictures by stating the relations they stand in to other things, or that there is a legitimate relational use of "is a picture of." But the relation between a person and their portrait is better understood in terms of a use to which a picture of a person (a person-representing-picture, in Goodman's terms) can be put, or a purpose it can have. This purpose can be contrasted with others, such as the illustration of a dictionary definition, or the illustration of a fictional world.[46] These purposes all relate pictures to things – to specific individuals, to words (and perhaps the concepts they express), to fictions. The kind of characterization of depiction that I turn to in this chapter is also a relational one, and corresponds to a distinct use of the locution "is a picture of." But it is not distinguished by a particular purpose that depiction can be put to, but rather a process by which pictures can be produced.

The process I have in mind is, roughly, that of making a depiction on the basis of one's observation of a specific object. The resulting picture stands in a certain relation to the object observed, and it can be considered a picture *of* that object, in a certain sense. I'll use the term "drawing from life" to refer to the practice of making pictures like this, though both "drawing" and "life" must be understood in an extended sense – this process need not involve the specific tools and techniques we associate with drawing rather than any other kind of picture-making, and it need not involve the observation of living things, though

[46] It would be interesting to compare this thought and the ideas behind it to Kendall Walton's view of pictures as "props in games of make-believe" (Walton 1973, 2003). Though what I have said is not compatible with Walton's view that the idea of a prop can explain what pictures *are*, his discussion nonetheless provides resources for thinking about what we use pictures to *do*. It would also be interesting to consider how it aligns with Noë's discussions of pictures as substitutes and as tools for putting things on display (Noë 2012, 2015). Much of what I have said is in the spirit of his idea that depiction is a matter of showing, but we disagree on the details of what pictures put on display and how.

people and areas of natural wilderness are canonical examples.[47] My general aim is to understand the philosophical interest of the fact that pictures can be produced in this way.

First, I'll investigate a question from Wollheim about the epistemological significance of drawing from life, and will use it to get clearer about how pictures teach us about the looks of particular things and the nature of the experience of seeing things like them. I'll do so by investigating in some depth the way we learn how to draw from life, and the practice's relationship to other activities and skills. Finally, I'll discuss the relationship between drawing from life and realism. Realism in depiction is often thought to be a product of investigation into the way things in the world really look. I'll suggest a way of understanding that without conceiving of it as a matter of accuracy or detail of any particular kind. I'll propose that we conceive of realism instead as kind of depiction that captures the individuality of visual objects and scenes. The idea of drawing from life, as an activity that involves sustained visual contemplation of the specific features of actual visible objects, gives us a grip on this feature of depiction, though realistic pictures can also be based on the imagination.

1. The new distinction

Drawing from life, and the difference between it and drawing from the imagination, is very rarely scrutinized in philosophical discussions of depiction. This might be because philosophers tend to take portraiture as the paradigm case when considering the relations that things can stand in as pictures. When we think of a portrait, we tend to think of a person "sitting" for that portrait – allowing an artist to observe them while they work on the picture. The image of this situation involves two different relations between the person and the picture, but the distinctive role that the person in this situation plays as an object of observation is likely to be overlooked, because the point of the image is to illustrate the concept of portraiture. Given this, it will be useful to consider the relationship between portraiture and drawing from life in some depth.

It is often claimed that the picture-maker's intention plays an essential role in determining who or what is portrayed by a picture (Wollheim 1980, Hyman 2012, Peacocke 1987). But this intention must be understood as necessary, not sufficient, and the content of the intention implicates requirements for the experience of the viewer. The intention of a portrait-maker is that their picture be *seen by viewers as a portrait of so-and-so*. This requires that the subject be *known* to viewers of the portrait. With realistic portraits of well-known people, this can be satisfied in a glance at the picture. With less realistic portraits and less well-known people, the picture must be presented in a more specific context, in conjunction with some indication of the subject. Often this is provided by the title or a caption. For example, without the title of Picasso's *Portrait of Gertrude Stein*, viewers at the time when it was made would not have taken it to be a portrait Gertrude Stein just by looking at it, even though they would have been familiar with the writer, and with her face. Most viewers of Alice Neel's *Ian and Mary* are introduced to the couple it portrays only by the title of the painting. They take the title to give the names of the people it represents because they know it belongs to a series of portraits of her acquaintances in Harlem in the 1970s. In contrast, the portrait of Mao Zedong displayed at Tiananmen Gate requires no further stage setting.

[47] For suggestions about the relationship between the difference between painting and drawing and the activity of drawing from life, see Berger 2016.

Drawing from life corresponds to a different kind of relation that an individual can bear to a picture though both can be expressed by the locution "is a picture of." This alternative relation is an etiological one, not a representational one. A picture can be said to be of a particular object when the picture is made by using the object in a certain way – when it is made based on simultaneous visual observation of that object. A picture might be made by using an object in this way, and yet not represent that object. A picture-maker might use an object in this way, without intending that the picture be seen as referring to that object (or any object). This use of "is (a picture) of" is closer to "is made of" – but it is also importantly different from many uses of "is made of." The object is obviously no part of the material that makes up the picture – that role is occupied by the picture's support (e.g. paper, canvas) and the medium used to mark it (e.g. ink, paint). Rather than a part of the product, the object is a part of the picture-making *process* – but its role is very different from that of a tool like a paintbrush. It plays this role in being there for the picture-maker to look at while they work on the picture, to guide the process of marking the support. Part of my aim in the subsequent sections of this chapter is to get clearer on what this role is, and how a picture-maker is guided by an object used in this way.

The object that guides the process of making a picture in this way *may* also be represented (in a relational sense) by the picture produced. As noted above, this is often the case for portraits: people often sit for an artist, and have their portrait painted while the artist looks at them. But though portraits may be based on visual observation of their subjects, they need not be. Gertrude Stein sat for Picasso 90 times, but in the end he scraped away all the work he did on her face during these sessions, and painted her face without looking at it. Hyman claims that it is possible for a fictional or mythological figure to be the referent of a picture (Hyman 2012, 135). Setting aside his use of the notion of reference, if it is correct that whatever the relation between a portrait and its subject is, it can be instantiated by fictional and mythical entities, then there can be portraits of Zeus and the Wizard of Oz. But Zeus and the Wizard of Oz could not sit for these portraits.

Conversely, a person can observe an object and base a picture on it without creating a picture that refers to that object. The notion of a *model* demonstrates this. A model poses in front of a picture-maker, who sketches the model based on how they look. But the sketch is not intended to be seen by anyone as referring to the particular person who posed.[48] The fact that we talk about models as "posing" rather than "sitting" for pictures marks the distinction between portraiture and drawing from life. Considerations about the use of models in picture-making make vivid that not all drawing from life results in a portrait. Painters often use studies of various different people and objects in conjunction in a single painting. One model may pose for a depiction of a man's robed torso, another may lend his head, another his hand. The finished product does not represent the body parts of these three different individuals. It may represent a different individual, such as Socrates. Or it may not represent any particular individual at all.

The role of a model in the process of drawing from life is comparable in certain ways to that of a photographed object in the process of making a photograph. In both cases, the object has a hand, so to speak, in determining the way that a surface is marked. And in both cases, the surface is marked to compose a picture of a certain kind of thing – a kind that applies to the object that determined the way the surface was marked. The manner in which the object determines the marking is different: in the photographic case, it is by structuring

[48] A drawing teacher might refer to the model in explaining where a sketch succeeds or fails, but that does not imply that the sketch is a portrait.

the light that a photosensitive material is exposed to; in the case of drawing, it is by affecting what the picture-maker sees. But in both cases, the object is part of the process of making the picture, and need not be represented by the finished product. In advertizing, for example, photographs of people and things are often used not to represent to the individuals photographed, but to represent a scene or activity of a certain kind. Photographs on the covers of novels often represent the characters in the novel, rather than the people photographed.

2. Representations as criteria for visual experience

Having homed in on this new sense of "is a picture of," we can begin to consider its philosophical interest. It is, in the first place, a matter of a picture's etiology. But this etiology has epistemological consequences. When drawing from life, the way the picture turns out is determined in part by what the object is like, upon observation. So when a picture is based on observation of an object, it is poised to convey knowledge, of some sort, about the object. So the phenomenon of drawing from life points us to an epistemic role that pictures can play.

In the essay "On Drawing an Object," Wollheim defends an idea about the specific kind of knowledge that pictures of this sort can capture and convey: that it is knowledge of how things look. He finds this idea in Wittgenstein's remark that the "criterion of the visual experience" is the "representation of what is seen" (Wollheim 1974, 3).[49] Wollheim doesn't use the phrase "drawing from life," but it is clear that this is the kind of picture he has in mind when he talks about "representation of what is seen."[50] The role of a "criterion" in this context is to provide a direct answer to the question of what a person's visual experience was like, or what kind of visual experience they had on a particular occasion. Wollheim makes clear that answers do not refer to candidate visible objects that the person might have seen, but rather to the *way* they saw a given object – or in other words, how some object looked to them. Thus, from a drawing of an object, we can learn how it looked to a person in certain conditions, and more generally, how it looks to people in those conditions.

Wollheim's primary concern in the essay is to defend this idea against the charge that it is absurd. By identifying certain bases for calling it absurd as tempting but mistaken assumptions about how the mind works, he clears the ground for a better understanding of how vision and depiction are related. The charge of absurdity arises from consideration about what the claim implies for the person whose visual experience is in question: can *they* learn from a drawing what *their own* visual experience is like, such that the drawing is a criterion for *them*? This is thought to be impossible, and so if the more general claim implies it, it can be rejected by *reductio*. Wollheim maintains that it is not impossible – it only seems to be so when faulty assumptions are in play.

The assumptions Wollheim identifies are related to a conviction that I have been inveighing against in other chapters: that seeing an object entails knowing how it looks. In the terminology of Wollheim's essay, this can be expressed as the idea that having a visual experience entails knowing what it is like. This makes it mysterious what a person could learn, from a drawing or by any other means, about the look of an object they have seen. It also makes it mysterious how a drawing could serve for *others* as the criterion, the definitive answer, for what a person's visual experience was like. Seeing the object they saw, in the

[49] The remark is §145 of *Philosophy of Psychology* (Wittgenstein 2009, 208).

[50] It is not entirely clear whether either Wittgenstein or Wollheim considers the criterion to be a *portrait* of what is seen.

conditions in which they saw it, would seem to be the ultimate step in investigating the nature of their visual experience.

Some of what Wollheim says to question these assumptions is quite convincing, and helps considerably to make room for an understanding of how knowledge of things' looks is hard-won. But ultimately Wollheim misconstrues the distance between the experience of seeing an object and the knowledge of how that object looks. In defending the claim that a person can learn from their own drawing how a thing looks to them, he assimilates two importantly different routes to the discovery of how an object looks: that of the person who draws the object, and that of the person who simply beholds the finished drawing. Recognizing and understanding the difference between these routes is essential to a coherent view of the epistemic role of pictures and their relationship to visual experience.

Let me first present the part of Wollheim's discussion that I think we should accept. This has to do with the way in which one's experience of seeing an object is the basis for any determination of how it looks. Wollheim imagines an opponent alleging that in order to use a picture as a criterion for the look of an object, one must, paradoxically, already know how it looks. Only then can one know that the picture is the kind that allegedly counts as a criterion – the kind that shows how the object looks (Wollheim 1974, 7). A picture can be used to convey knowledge about how an object looks to those who haven't seen the object. But in order for it to be used in this way, the person who makes it and presents it to others must have checked the picture against what they already know about the look of the object. So a picture can't be the ultimate source of this knowledge.

In replying, Wollheim admits that the *visual experience* of the object must be independent of the production of the picture, and must be the basis of it, in a certain sense. But, he points out, that is not the same as the picture being based on *knowledge of* the visual experience. The idea that a picture can be a criterion for how an object looks to a person requires that there *be* some way that the object looks to that person. And the production of the kind of picture that can serve as a criterion requires that the picture-maker have some experience of seeing the object. But the picture-maker need not know in advance how a thing looked to one, in order to assess whether a drawing adequately represents it (Wollheim 1974, 7).

Completing the picture will, however, require judging that the work that has been done is adequate to show how the object looks. But, Wollheim explains, this judgment need not involve appraising the picture in light of independent knowledge of how the object looks. The fact that the picture is adequate and the fact of how its object looks can be discovered simultaneously, in a single experience of looking at the picture. The experience of seeing the object is required: it puts one in a position to consider whether a picture is successful in capturing what it is like. But again, having seen the object, and having knowledge of what that experience is like, do not come to the same thing. What is necessarily prior to the depiction is simply the visual experience itself, not the stock of knowledge of its nature that a finished picture is able to confer (Wollheim 1974, 9).

Wollheim connects resistance to these points to the idea that the experience of seeing something involves having an image somehow in mind, an image which may linger as we then try to reproduce it with pencil and paper (Wollheim 1974, 9-10). The judgment that a picture is correct, then, is a comparison of two things whose natures are apprehended independently. Apprehension of the mental image is, of course, prior. This way of thinking encourages the conviction, that having a visual experience with some particular phenomenology comes along with *knowledge* of its phenomenology, insofar as apprehending the mental image is understood as having knowledge of the appearance what one is looking

at. Wollheim does not dispute the coherence of the analysis of seeing as the occurrence of pictures in the mind. Instead, he argues that there is no reason to think it is true. Specifically, the practice of drawing from life does not provide any: to explain it, we do not need to posit that there are images in our head for us to copy. We can make sense of the drawing process as a process of finding out what our visual experience is like, as something that our visual experience of an object merely puts us in a position to do.

These ways of responding to the charge of absurdity seem sound to me. I am convinced that we can make sense of drawing from life as an activity that expands our knowledge of our own visual experience, rather than simply recapitulating it. But there remains the question of *how* to make sense of it as affording this possibility. I think Wollheim's further claims about the activity get in the way of a proper understanding of this. I'll turn to those now.

Wollheim addresses the suggestion that a person cannot possibly learn what their visual experience is like from a drawing they make, because making the drawing must be an intentional action, and we cannot learn about our own intentional actions by observation. If the drawing is of the kind that conveys knowledge of the nature of one's visual experience, then it must be made with the aim of showing what the object looks like. But then the person will know what their drawing shows simply in virtue of being the person who makes it, prior to any observation of what they make. So the picture may convey to other people how the object looks, but it cannot be the picture-maker's route to that knowledge (Wollheim 1974,10-11).

Wollheim responds by claiming that there are in fact certain things we do intentionally that we find out about by observing our actions. Put more precisely, the point is about descriptions that can characterize our behavior as intentional, and the knowledge that these descriptions apply. In certain cases, a description can characterize a behavior as something we do intentionally, despite the fact that we can learn from observation that it applies. To illustrate the point, Wollheim describes a person calculating an arithmetic series, saying the sequence of numbers out loud, as he calculates them. The person might not know when he begins that the seventh number in the series is, say, 52. Thus he might discover what the seventh number in the series is upon reaching the seventh place, by hearing which number he says. Despite that, he says "52" intentionally (Wollheim 1974, 17).

Wollheim thinks that the intention to carry out the arithmetic series is analogous to the intention to draw an object one sees, and that the discovery of what the numbers in the series are is analogous to the discovery of what the object looks like, as represented by the picture one makes. Both discoveries, he thinks, are made by observation of one's own action, or the product of one's intention. He suggests that this epistemic relation holds generally for actions with a certain structure, such as applying a rule in various different instances (Wollheim 1974, 17).

The analogy with constructing an arithmetical series does not serve to establish Wollheim's claim, however. It is true that a person who decides to construct an arithmetic series need not know any of the numbers it contains before he gets started. It would be misleading to say that, at the outset, his intention is to say "52" at a certain point in constructing the series. And it is true that despite this, when he arrives at a certain point in constructing the series, he says "52" intentionally. But it is not true that he learns the numbers in the series *by listening to himself saying them*. Rather, he learns them *as he calculates*, in intentionally constructing the series. The construction of the series can be thought of as an inquiry into the numbers it contains, and the numbers are discovered in carrying out the inquiry, not by an observation of the inquiry. If the drawing case is analogous, then we

should not think of it as a violation of the principle that one does not learn about one's intentional action by observation. The picture-maker learns how the object looks in carrying out the drawing project, not in observing what she carries out.

One might think that this is all irrelevant to the issue of whether a person can learn by observing their own drawing what their visual experience is like. The objection notes the epistemic status of observation of one's own *actions*, and a drawing is not an action, but the result of one. So why does Wollheim concentrate on the result? I think he treats the objection this way because of how he conceives of the claim he wants to defend. He is interested in pointing out the epistemic distance between having a visual experience, and having a representation of it. He wants to do this by defending the claim that a picture-maker is no better off, epistemically, than the other viewers they show their picture to, simply in virtue of having *had* the experience to be conveyed by the picture. And he thinks that to maintain this claim, he must maintain that the picture-maker learns how the pictured object looks in exactly the same way as the other viewers, by looking at a finished drawing of it. He overlooks the possibility of learning how an object looks *by drawing* it, rather than by looking at the finished product. This still posits distance between having the visual experience, and knowing what the object seen looks like. But it allows that this distance can be traversed in different ways: by making a picture, or by being shown one.

To make room for the first option is *not* to identify a way of learning how an object looks that is manual, rather than visual. Drawing from life is a manual activity, but it is a visual one as well. Wollheim himself acknowledges that we use our eyes to draw, looking at the surface as we mark it, making sure that the marks are "on track" to realize the aim of our drawing (Wollheim 1974, 13). But he takes pains to show that the look of the object can be discovered through looking that is not part of the drawing process, but rather counts as observation of the finished product. His point is that this *can* happen, not that it is *always* how a picture-maker discovers the look of the object she draws. But he suggests that when this does not happen, the picture-maker *does* know how the object looks before undertaking to draw it. So it seems that for him, there is no discussion to be had about how a person can come to see how something looks in the process of drawing it.

The looking involved in drawing from life is not only navigational – a means of making sure one is putting certain marks in certain places. It is also a means of contemplating the things one sees. And what one sees while drawing from life is not only one's drawing as it proceeds. It crucially involves looking at something *in addition* to the drawing: the object being drawn. Having got that object in view, drawing it involves further visual scrutiny. So there is a sense in which the picture-maker learns about the nature of her visual experience by observation. But the observation that provides this new knowledge is not observation of her finished drawing, nor is it of the activity of drawing that she engages in. The observation is *part of* that activity, and it comprises sustained observation of the object she sees, as well as of her drawing surface and the marks she is making on it. Wollheim would be right to say that this knowledge of visual experience is won by observation of external objects, as opposed to introspection. And he would also be right to say that observation of the object need not provide an independent criterion on which judgments about how to draw it are based. The way the picture-maker observes an object may be dependent on the simultaneous observation of her drawing surface as she modifies it, such that the use of the medium is integral to obtaining this knowledge of the object. But that should not lead us to say that she obtains this knowledge by observing her own activity

of drawing, or that it is available to her only after she finishes the picture and looks at it again, as a record of that process.[51]

Perhaps Wollheim recognizes this obliquely, in addressing a rather vague worry that a picture cannot convey what an object looks like, because a picture and the object it depicts look so different. In responding, Wollheim discusses the role of line and contour in drawing, and his remarks could teach us something about the way that these features of pictures are used to make clear how the object looks – clearer than it usually is to someone who looks at the actual object itself. But these remarks are focused on the perspective of a viewer of the finished product, someone who is shown how the object looks by the picture. And (as I have argued) this is not the perspective of the picture-maker. We need an explication of how the process of *making* marks can constitute an inquiry into the look of an object.

In her essay "Drawing From Life," Antonia Phillips makes progress on this front. Her aim is to define what it is for something to count as a drawing of an object from life. Having established that the object must be seen by the picture-maker while making the picture, and must be involved in some way in the process of producing it, Phillips says how the object must be involved. She says that the picture-maker must have a specific skill, and must produce the picture by deploying that skill on the object. The account of a certain kind of representation goes by way of an account of a certain kind of skill. In the next section, I'll present Phillips's view of this skill.

3. The skill of representing what one sees

Phillips's aim is to define what it is for something to be a drawing from life. What is distinctive of drawing from life is the role that a particular visible object plays in the depictive process. The process is a physical one, but, she thinks, it cannot be defined in mechanical terms. Instead, we can think of a *skill* as "lying at the heart of the activity" (Phillips 1992, 322). Phillips notes that this approach will work only if it is possible to identify the skill without referring to the aim of drawing from life itself. To do this, she looks to the history of artistic training, specifically to a traditional course of study that began at the end of the fourteenth century in Western Europe.

The traditional course has three stages, culminating in mastery of drawing from life. In the first stage, students practice copying other pictures. In the second, they are presented with sculptures and plaster casts to depict. In the final stage, they make pictures based on live models and landscapes. Here, they are actually practicing the activity of drawing from life, having been prepared to do so by the different activities of the first two stages.

Phillips asks what the "special contribution" of the final stage amounts to, thinking that this is the key to understanding the skill of drawing from life. Her answer is that the activity practiced in the final stage requires students to make a set of choices, distinguishing certain elements of the visible world from the rest. Any object or scene admits of more features than any picture could capture, and so a person who aims to draw something from life must "select elements in what he sees for transcription as elements of his picture" (Phillips 1992, 324). The prior stages of training prepare students to meet this challenge,

[51] Berger writes about the experience of finishing a painting in "Painting and Time." He says there that it is not a matter of achieving a correspondence between the painting and "something already existing – like the second shoe of a pair," but rather a matter of "when the *foreseen* ideal moment of it being looked at is filled as the painter feels or calculates it should be filled" (Berger 1986, 206). This is related to what I am saying here, namely that the painter (or picture-maker) need not inhabit the role of the spectator of the picture in order to know that it is done. The painter can *foresee* the experience someone else will have when they are shown the picture, and know that it will be as they intend. They do not need to have that experience themselves.

presenting them with tasks that require related but less difficult problems of selection. At the first stage, when copying pictures, students do not have to determine compositions for themselves: the questions of what kinds of marks to use and where to place them have been answered for them. At the second stage, the student has to make decisions about how and where to mark a pictorial support, but in relation to "forms already wrought into art" (Phillips 1992, 324). So the elements that the student must "transcribe" are more or less chosen already, though there are decisions to be made about how to accomplish the transcription.

Phillips acknowledges that there may not be a "firm line" between the activities of the second and third stages. In both cases, one is making a picture of what one sees, whereas in the first, one is copying a picture one sees – producing a picture of something that has already been depicted, not a picture *of a picture*. But Phillips thinks it is important that the third stage involves drawing things that are "untouched by pictorial and artistic conventions" (Phillips 1992, 324). She thinks that because of this difference, the third stage instills a new skill "with a generalizable nature," one which enables the student to "make a picture out of *anything* he sees" (Phillips 1992, 325).

But even with the admission that the division might be fuzzy, this way of contrasting the second and third stages is not convincing. We can't ignore the fact that the basic activity is the same: at both stages, students are asked to draw what they see. In the third stage, the student is expected to develop a more general ability, to successfully engage in an activity in a broader range of circumstances, including ones that are more difficult. The increased difficulty of the exercises introduced at the third stage means that the student's skill must develop to meet new challenges, but it develops in degree – the skill required is the same in kind. Whereas Phillips's conception of the relationship between the last two stages is somewhat like that of biking with training wheels and biking without them, it is actually more like that of biking on a road that is smooth, straight, and flat, and biking on more treacherous terrain.

But the fact that the skill of drawing from life is at play in stages two and three does not mean that Phillips's general strategy is misguided. Reflecting on the progression of the traditional curriculum can still help to clarify the skill of drawing from life. In fact, it means that the curriculum provides more help in this connection than it otherwise would. It gives us two things to consider: the difference between the activities of the first stage and the final two, and the fact that sculptures and inanimate artifacts typically provide easier practice of the same skill.

At all stages, the student's activity is both a manual and a visual one. The gestures of mark-making by hand require dexterity that the student must acquire. It takes practice to understand the behavior of any materials one might be using, and to learn how to handle them. But the mark-making activity is always also a visual one. The student practices modifying her handling of the medium as she goes, based on the effect she sees herself having on the drawing surface. She learns, for instance, to reduce pressure on the marking tool while pulling it across the surface, to create a line that tapers at a certain rate. Once she is fairly familiar with the behavior of the tool, she had developed robust expectations of the effect it will have, and might be able to mark the surface in a specific way without looking at it. But her eyes will always have to come back to the surface eventually to check her work. Even when things go exactly according to plan, vision is her way of keeping the drawing on track. The student must develop the ability to tell, by looking at her drawing, whether it is on track or not. She must be able to see the relationships between different marks, and the shapes they contribute to.

In all three stages, the student exercises visual skill not only in looking at the surface she draws on, but also in looking at something else. In the first stage, she is given pictures to copy. Phillips says that this stage is easier than the others because the choices that go into making a drawing – selecting which features to represent things as having, and determining the kind of marks to use and the placement of them – are not required of the student, but are done for her. But this obscures the nature of the activity students undertake at this stage, and the distinctive difficulties of undertaking it. In this stage students produce images, but they do not *draw* things; rather, they *copy* drawings. Their task is to reproduce the same arrangement of marks that appears on the surface they are given. This activity provides fairly focused practice in handling of the medium, and develops their dexterity in mark-making. But it is also a visual challenge. It is difficult to produce a copy of something just by looking at it; it requires exercising the visual skill of judging the relationships between marks on a surface on both the picture she is given and the one she is producing. Reproducing something is easier if you can trace it, or use measuring tools on the original and the copy. But in this activity, students are deprived of either of those methods, and must use only their eyes to ascertain and reproduce the relevant relationships between marks. There are plenty of choices, or determinations, that a student has to make at this stage, and they are not easy to make. Pictures are particularly difficult sets of marks to reproduce; an abstract arrangement is often much easier. Students are often invited to turn the picture they are given upside down, so that it looks more like an abstract or haphazard set of marks than one with pictorial significance. This makes it much easier to do well at the activity.

The activity of copying is a challenge, and requires visual skill. The skill that the student develops at the first stage is involved in the activity of drawing from life, but it is not identical with it. The skill required in the later stages can be distinguished from it, and the fundamental difference between the two activities can be understood in light of it. One key to the distinction is how the success of the result of each activity is judged. The other is why the skill behind drawing can be built up by moving from sculptures and still lifes to live models and natural scenes.

Whether a student succeeds at copying a drawing can be determined by mechanical means. One only needs to check whether the markings of the two drawings *match*, and their congruity can be determined simply by measuring them. But the success of the activity of the subsequent stages cannot be determined by any such mechanical method. It can only be judged by a sufficiently skilled viewer, who can look at the drawing and the scene or object it is based on.

The observed object plays a different role in the judgment in either case. We can make sense of this by realizing that it plays a correspondingly different role in guiding the drawing as it proceeds, and by pinpointing what role it plays in the later stages. In stage one, the student is concerned with what the things she is looking at are like (how the marks on a surface are shaped and placed in relation to each other), but not necessarily what they *look* like. Having ascertained that visually (and through manipulating her medium), she has determined how her own marks should be made. In stage two and onwards, it is essentially how the observed object *looks* that determines how she should proceed. She must then find a way of *showing* how the object or scene looks using her medium, a way of marking the surface to make this way of looking visible.

This analysis of the activity of drawing makes sense of the fact that its success cannot be determined purely mechanically. It can only be judged by a sufficiently skilled viewer, who can look at the drawing and the scene or object it is based on, and determine whether and how well the one shows how the other looks.

It also makes sense of the fact that it is generally easier to draw sculptures and purposeful arrangements of inanimate artifacts than real people and scenes that are haphazard collections of natural and artificial elements. Sculptures are objects like drawings, in that they are designed to show how things look, using a different medium. The vastness of the difference in media means that they cannot be copied. But the fact that they are representational media means that the way they look is particularly salient. Compared to the real article, a sculpture of something will often present the visual relationships between its parts explicitly, making them prominent. Thus it is relatively easy for a student to discern what she needs to see to draw them. It may do so by having a simplified shape, an approximation of more complex topology. It may do so by having a uniform texture, eliminating the contribution that variation in texture makes to visual appearance, revealing the distinctive role that shape plays in how things look. Inanimate artifacts typically have simpler shapes than organic objects, and they can be arranged in ways that make their visual appearance simpler, and easier to see and convey.

Only once a student has had success at the third stage, at a higher lever of difficulty, do we say that they have mastered the activity of drawing from life. But this is not because the ability comes into play only in the third stage, and not in the second. Rather, it is because the second stage involves a very limited range of circumstances in which the ability can be exercised. Having succeeded at the second stage, we can only attribute the ability in a qualified way: the student can draw from life *if* what they set out to draw is sufficiently simple. With success at the third stage, we can attribute the ability without any such qualification; the set of circumstances in which the student is equipped for success is entirely general. The ability of the intermediate student is more limited than the advanced one, but the particular ability they have is one and the same; they are practicing its exercise as soon as they are told to look and *draw*, rather than to look and copy.[52] From that point onwards, they are developing the skill of seeing and showing how things look.

To conclude this section, I want to connect this analysis back up with the question of how a picture can serve as the criterion for a visual experience. The claim was that if someone makes a picture of an object by drawing it from life, then the picture shows how the object looked to the person who made it. The idea that drawing from life is based on the skill of seeing how things look, and manipulating a medium to make marks that show the looks of things, might seem to align quite well with this. This alignment was part of my aim, and I hope to have illuminated how pictures can play this role by examining the skill behind making them and comparing it to others in the vicinity. But one thing I've said may raise a worry about holding on to Wittgenstein's claim that pictures can serve as criteria for experiences. I've said that the success of a drawing of an object is judged by someone who can look at the drawing and the object drawn, and I've implied that the object is *used* to judge the drawing (in a different way than a drawing is used to judge a copy, in the first stage). One might worry that this means that seeing the object is the criterion for the success of the drawing, and so a drawing can't be the criterion for the visual experience of the object.

---

[52] The use of photographic references might count as a second-stage exercise. When one's medium differs considerably from what one is looking at, one may create a drawing by way of something more like translation than copying. However, it could be said to differ importantly from the activity of drawing sculptures in that the result in one case is a drawing of the object looked at (a sculpture), while in the other case the result is a drawing of what is represented by the object looked at (the photographed object). If Walton's claim about the transparency of photographs is correct, however, this difference would disappear (Walton 1984). It seems, in any case, that the distinction between the first and second stages may have its own kind of fuzziness as well.

We must distinguish between the context of training to develop the ability to draw from life, and the context of exercising that ability once it is established. I've been talking about the former, while Wollheim's discussion is focused on the latter. He notes that for a picture to count as a criterion for a person's visual experience, we must take it to be the product of that person's ability to draw things from life, applied to the object that they saw on the occasion in question. But we can know that the picture fits the bill without judging it against the object. We need to know that the person has the ability to draw from life. But that general fact about the person can be established without reference to the object that they drew on this occasion. To establish their ability, it may be necessary to compare some of their drawings with the objects they've attempted to draw. But this need not be *reestablished* for every new drawing.

But a further clarification is in order about the judgment that we make of someone's ability to draw from life. Though it involves comparing their drawing with the things they have drawn, it need not involve *first* ascertaining how their models look, and *then* checking their drawing against this. Their drawing might reveal aspects of a model's appearance, which we see in looking at the model only after seeing the drawing, enabling us to determine that it is a successful drawing from life. The object itself does not determine any one way of representing it, and in judging whether it has been successfully represented, one need not be an authority on any set of rules for drawing objects.

4. Realism

I've been discussing the use of drawing from life as a way of conveying a matter of *personal experience*: how a particular thing looks to an individual on a certain occasion. But it is also connected to a broader concern: the project of depicting the world *realistically*. They are not the same project – there are realistic pictures that are not drawn from life. But they are interestingly related. In this final section, I'll try to shed light on what realism is by considering how it relates to the activity of drawing from life.

Realism is a feature we attribute to some pictures and not others, and one that we take to come in degrees. While we can judge pictures to be highly realistic in comparison to others, it is not clear whether there is an absolute maximum of realism. Realism is often contrasted with stylization, and thought of as "natural" representation – "naturalism" is sometimes used as a synonym for "realism." But it is important to keep in mind that the phenomenon of stylization needs explication just as much as that of realism of naturalism. It is also important to keep in mind that stylization is not the same as *having a style*, lest we make the mistake of considering realism to be the absence of style.[53]

Realism is not the absence of style, nor is it a particular style. Some philosophers take realism to characterize styles as well as individual pictures. Here I will focus on what it is for an individual picture to be realistic. One of the views I will consider explains this as a function of a picture's style, but in doing so it does not tell us what it is for a style to be realistic or for one style to be more realistic than another. Regardless of how we handle the question of attributing realism to pictorial styles, realism as applied to individual pictures can be considered as a *manner* of depicting – a matter of how rather than what a picture represents.[54]

---

[53] This is one of the central themes of Gombrich's *Art and Illusion* (Gombrich 1960).

[54] There is a use of "realism" that refers to an artistic movement characterized by the choice of unromantic subject matter, but I will consider that to be a separate use.

It is difficult to get much further than these structural features of realism without venturing into philosophical controversy. The concept is often unpacked in terms of notions like fidelity and truth, but conflicting stories abound as to what a realistic picture must be faithful to (the retina? the laws of optics? the behavior of sense-data?). It also common to find realism unpacked in terms of the notion of illusion – and just as hard to find consensus on how to spell out what the relevant illusion is of, or under what circumstances a realistic picture will produce it.

To get a grip on the notion of realism, some point to the history of pictorial art and the way pictorial techniques and styles have changed over time, and ask us to see this chronology as gradual progression towards an ideal.[55] If we can characterize this ideal in other terms, we will have shed light on what realism is. This approach is not uncontroversial, in that it assumes that the development of realism amounts to progress, and this premise might require defense. It also brings out how abstract and elusive the criteria for realism must be, since the path of progress seems to be a branching one: there are very different styles that all seem to be realistic. The diversity of these ways of depicting makes it difficult to identify a unified ideal that they all achieve.

It is clear that our ways of picking out realism as a phenomenon raise difficult philosophical questions almost immediately. I do not take that to signify a failure to pick out the phenomenon, or as evidence that the concept is incoherent. Instead, I take it as an indication that the concept of realism may not be entirely pre-theoretical. If it gains most of its substance in the context of reflecting on what depiction is and what it accomplishes, then there will be relatively little to say about it in advance of any philosophical investigation. I'll proceed from here as though the target for understanding has been located, even if its contours are not entirely perspicuous.

I'll begin with Goodman's treatment of realism, which rejects the idea that it is a matter of illusion, or fidelity, or the achievement of any particular objective ideal. Goodman begins with the analysis of realism as "deception" – the claim that a picture is realistic "just to the extent that it is a successful illusion, leading the viewer to suppose that it is, or has the characteristics of, what it represents" (Goodman 1976, 34). This is better, he thinks, than the view that a picture is realistic to the extent that it *actually has* the characteristics of what it represents. But he points out that it conflicts with basic facts about our use of pictures. For one thing, if a picture were guaranteed to deceive any viewer, it would not thereby count as depicting something entirely realistically – it would be impossible to see it as depicting at all (Goodman 1976, 34). Secondly, we are rarely ever deceived by the pictures we look at. If we "enlist mischief" in setting up special viewing conditions, we can get a picture to deceive someone. But this is not proof of the picture's realism. Even an extremely unrealistic picture can be deceptive if its viewing conditions are manipulated (Goodman 1976, 35).

Goodman thinks the failure of illusion to serve as a measure of realism leads to the idea that it has to do with the information that a picture conveys. On this view, "the most realistic picture is the one that provides the greatest amount of pertinent information" (Goodman 1976, 35). He rejects this claim on the basis of an alleged counterexample: two pictures, one of which is realistic, and the other of which is "just like the first except that the perspective is reversed and the color is replaced by its complementary." He claims that "appropriately interpreted," the second picture provides just as much information as the first. But it is clearly not realistic (Goodman 1976, 35). The same example is used to prove

[55] This is part of the way that Gombrich approaches realism in Gombrich 1960.

that accuracy, or fidelity, is neither equivalent to nor sufficient for realism: the pictures are equally faithful or correct, but only the first is realistic (Goodman 1976, 36).[56]

The relevant difference between the two pictures, Goodman thinks, is that the appropriate interpretation of the second one will not be as ready to hand. He identifies this as the "touchstone of realism": it is not a measure of "quantity of information" but "how easily it issues" (Goodman 1976, 36). According to Goodman, the latter depends in turn on how "stereotyped the mode of representation" of a picture is. It is relative to the viewer, and determined by which ways of making pictures they are more or less familiar with, given their cultural context: "Newer or older or alien systems are accounted artificial or unskilled" (Goodman 1976, 37).

But this doesn't seem to square with the fact that viewers can be relatively well acquainted with a wide range of pictorial traditions, produced in cultures that are distant both geographically and culturally, and yet still consider certain pictures and styles more realistic than others. This seems to be the case for most 21st century viewers, and I'd bet that it was the case for viewers reading *Languages of Art* in 1973. There is certainly a history of conflating realism with artistry, and a history of refusing to recognize the achievements of foreign cultures. But we can distinguish between realism and artistic skill or merit, and we can approach the art of any culture with clearer eyes. When we do, the concept of realism still gets a grip. We can be well acquainted with a wide range of styles of depiction, and distinguish degrees of realism in pictures of various styles. So our judgments of realism do not seem to depend on the familiarity of a picture's style, and realism itself cannot be chalked up to the ease with which we can see a picture as depicting what it does.

The fact that Goodman's view is untenable does not that mean that any of the views he rejects is correct. But many philosophers have attempted to rescue the idea that realism is a matter of providing pertinent information. Abell suggests the motivation for this: a pretheoretical intuition that realism has a "special epistemic role" which must be acknowledged and explained by a theoretical account of it. As she sees it, more realistic pictures are "better sources of some sorts of knowledge" (Abell 2007, 2). Like Goodman, Abell denies that the realism of a picture amounts to either its accuracy, which she construes as the quality of information it conveys, or the level of detail it exhibits, construed as the quantity of information it conveys. But, she argues, it is nonetheless a function of the amount of pertinent information that it carries. Goodman's alleged counterexample proves to be irrelevant if we recognize that there are "constraints on the information that determines realism" (Abell 2007, 11).

One of these constraints is on the way in which a picture provides information. Realism is a measure of "*depictive* information: information that pictures provide because of their depictive content." This reveals the crucial difference between the two pictures Goodman describes. It is not the ease with which their information can be accessed, but rather the way in which it is carried. Both provide accurate information about the colors of the things they depict, in the sense that differences in the colors that compose the picture correspond to differences in the colors of the things it depicts. But only one of the pictures provides accurate information about these things' colors *by depicting* them (Abell 2007, 12). The other constraint is on the type of information a picture can provide. The information

[56] This last point depends on the claim that these pictures could not only contain the same information, but contain it by *depicting* exactly the same thing (be classified as the same kinds of pictures). Goodman claims that there is no reason to deny that they could, though this would require that they appear in different contexts, where different "systems" of depiction are in play.

that determines realism is information about "how an object would look were one to see it" (Abell 2007, 13).

Having focused in on the information about things' looks, Abell proposes that realism is a matter of whether this subset of the information a picture conveys is "relevant." Abell works with a conception of relevance as a matter of whether information "connect[s] with viewers' existent assumptions to yield positive cognitive effects that warrant the processing effort required to obtain them" (Abell 2007, 11). Cognitive effects are changes in the viewer's assumptions and the relative salience of those assumptions. Such effects count as positive if "like true conclusions and unlike false conclusions, [they are] worth having" (Abell 2007, 12). Abell does not specify a way of gauging processing effort. But the general idea seems to be that realism has to do with the way pictures provide information about aspects of things' looks that are particularly important to us. Which aspects are important will vary from viewer to viewer. It may depend on aspects of their cultural context, including the pictorial traditions that they have been exposed to. Abell gives an example that depends on practical significance: she imagines a group of people "whose staple diet is a plant that is visually very similar to a poisonous plant." For them, "information that serves to distinguish the edible plant's appearance from that of the poisonous plant will be very relevant" (Abell 2007, 16). This gives us a sense of how and when relevance will vary, but it does not define relevance in terms of any further kind of concern. On Abell's view relevance may be the result of any number of factors; whatever its source may be in a particular case, it determines a picture's realism.

Before we can assess either Abell's or Goodman's view of realism, we must note an ambiguity in the way they are stated, and determine which of their possible meanings is intended. Abell and Goodman both draw our attention to the information a picture provides about *its object*. But as we have seen, that phrase is ambiguous: it might refer to a picture's referent, or it might serve to classify the picture as a depiction of a certain kind. The first reading makes the claim easier to understand – we're familiar with the idea of one object conveying information about another, and we can easily apply that idea to the pair of a picture and its referent.

Goodman does not address this explicitly. Abell does, however, and ends up articulating two different definitions of realism, corresponding to the two senses of "depicts." In the final analysis, a picture is realistic "qua picture of a particular" insofar as "the depictive information it provides about how the particular it depicts would look" is relevant (Abell 2007, 11). To explain what it is for a picture to depict realistically in the non-relational sense, Abell simply substitutes "object type" for "particular" to produce a new statement: a picture is realistic "qua picture of an object type," insofar as the "depictive information it provides about how an object of the type it depicts would look" is relevant (Abell 2007, 11).

Abell claims that a portrait can be realistic in both of these senses: it can provide relevant information about the look of the particular individual that it refers to, and it can provide relevant information about the looks of things of the type it is depicted as being. If it is realistic in the first sense, it is also realistic in the second. But the reverse does not hold. To illustrate, Abell imagines a portrait of a particular man ("your uncle") which depicts him as a thin and bald. If he is in fact "fat and hairy," then the portrait is not realistic "qua picture of a particular." But it may nonetheless be realistic "qua picture of an object type," if it provides relevant information about how a thin, bald man would look (Abell 2007, 13). Abell privileges "realism qua picture of an object type," however, by calling it "realism *simpliciter*." It is "the most basic form of realism," since it allows us to "compare the realism

of all pictures," regardless of whether they have referents. She thinks is also "stands to reason" that this form of realism is the basic one, since "our default way of interpreting pictures is as of object types," and we often don't know whether a picture portrays an particular, or have any other sources of information about that particular against which to judge whether the picture provides accurate information about it (Abell 2007, 13-14).

In applying the two definitions to the example above, however, it becomes clear that there are not two different senses of realism, but only one – there is no such thing as being "realistic qua picture of a particular." If a picture portrays a particular fat and hairy person as being thin and bald, it is therefore inaccurate, but not unrealistic, in any sense. What Abell calls "realism *simpliciter*" is simply realism. The definition of "realism qua picture of a particular" points us to a different phenomenon, namely accuracy.[57] Abell agrees elsewhere in the paper that realism and accuracy should not be equated (Abell 2007, 3).

Realism characterizes depiction only in the non-relational sense. But if we take that sense to be in play, the statement a realistic picture provides relevant information about the look of its object conveys a very different kind of thought. We no longer have a pair of things to consider, one carrying information about the other. Instead, we have only the picture, understood as an instance of a given type of picture. In that case, what would be meant by the claim that the picture conveys relevant information about the look of its object? It would be that the picture provides information about *what its object is* – that is, about *what type of picture it is*. But does it? Certainly, the relevant facts can be gleaned by looking at it – a picture's object, in this way, is a visible aspect of its surface. But does that constitute providing information? Do I provide information about my height, by being 5'8"? It might be difficult to answer the question of whether or how an object can provide information about itself. But regardless of how we answer it, it doesn't seem that this can determine its realism. If there were a way in which pictures inform us about what kind of pictures they are, it's not clear why we shouldn't expect that all pictures would do so equally.

This is almost certainly not what Abell has in mind, however.[58] When she says that a picture is realistic "qua picture of an object type" in virtue of the "the depictive information it provides about how an object of the type it depicts would look" (Abell 2007, 11). The information this seems to point to is not about the *picture's* type or classification, but rather about the type of object referred to in classifying it. Specifically, it is information about the way that objects of that type *would look*. Usually, the phrase "would look" is followed by a specification of a set of circumstances, e.g. "under fluorescent lighting" or "from a great height." But Abell intends the circumstances to be entirely general: she means to invoke how the object would look *if seen*. Keeping in mind that there is no *particular* object at issue here, we must understand the information that matters as being about how objects of a certain type look, in general.

I'm not sure whether it makes sense, ultimately, to think of the way things look in general as information about them. I do think pictures can provide knowledge of how things

---

[57] Abell objects to Schier (Schier 1986) for "construing what is intuitively a single phenomenon, pictorial realism, as in fact comprising two separate phenomena" (Abell 2007, 8). It is not clear why she doesn't see this as a problem for her own view, or at least something in need of defense. It is possible that she does not consider the phenomena she invokes to be so separate, since their statements differ only by the substitution of a single phrase. (I'll explain just below how different they really are.)

[58] It may be more or less what Goodman has in mind. He might think that even though picture's types can be determined by looking at them, they can be easier or harder to see as depicting what they do, for different viewers under different conditions. Perhaps these are the differences that Goodman has in mind as determinative of realism. But I have addressed the problem with this view already.

of various types tend to look, and that viewers often do learn this from them. But setting aside questions about the role that the concept of information should play in our epistemology, I want to consider the subject matter at hand, the way that things of a certain type tend to look. I want to suggest that this is not the essence of realism, and that its actual essence is more or less the opposite.

Abell's definition has it that realism is a matter of accurate generalization about the way things of a certain type look. But intuitively, realism is about specificity. When one is struck by the realism of a picture, one has the sense that the picture is the way it is not because that's how things of this type *typically* look, but because visible things are *individuals*, and the picture conveys that individuality. A realistic picture seems to capture the way that contingent features of things contribute to the structure of a visible scene. A realistic picture makes manifest the fact that every particular thing we see is distinct from all others, and all seeing takes place in some set of contingent circumstances rather than any others.

Here (at last) is the connection between realism and drawing from life. A picture need not be drawn from life to be realistic. But drawing from life is an activity in which specificity plays a central role. In drawing from life, one aims to create a picture of a certain type, but also to respond to what is specific to a particular instance of that type, and how it happens to appear. If successful, a drawing from life conveys the look of a particular object, whatever it happens to be. The drawing process is responsive to the object's specific and contingent way of appearing. The resulting picture displays a sense of the specificity of visible objects, which is derived from an actual object – it is achieved by attending to the way that object is, in contrast to the way objects like that typically look and are typically represented.

This quality of specificity is precisely what realistic pictures exhibit. They impress on us the fact that a viewer always sees individual things in specific circumstances. It is a simple fact, but seeing a realistic picture can heighten our appreciation of it in a distinctive way. Though an examination of drawing from life makes clear what this effect is, it need not be achieved directly through the observation of a particular object. It may instead be achieved by an acute imagination. The effect is to present the viewer with the specificity of visible objects in the abstract, not to convey the facts of any actual object's appearance. It is not a matter of depicting things accurately, nor easily, but vividly.

It is also not a matter of depicting things in detail. Detail is orthogonal to the specificity that constitutes realism, since there is no particular level of detail at which we see things. Whatever level of detail we see in a thing, the experience is still an encounter with a specific object as opposed to any other. Depictions of detail in the play of light or in subtle modulations of surfaces *can* strike us as realistic. Illumination is a highly variable aspect of our surroundings, and so it often contributes to a sense of the uniqueness of a moment of seeing. Variations in surface texture and mottling also tend to heighten our appreciation of the individuality of visible things, since they often differ between instances of a given type of thing. But these features are merely correlated with an appreciation of the specificity of an occasion of seeing, and are not constitutive of it. Conceiving of realism in terms of specificity does not involve identifying it with any particular set of techniques, and instead explains it as an inherently open-ended phenomenon. Likewise, training to draw from life, as I've described it, does not involve mastering a set of techniques, forming certain habits, or adopting a particular style. Rather, it involves practice in looking at things one finds in the world (or is given by a savvy teacher), and being guided by one's observation of them. The skill it instills can be described as a visual and manual one that enables those who have it to

see the looks of things they encounter in the world, and to see how markings on a surface can capture their look.

The activity of drawing from life affords discoveries about particular objects, but those discoveries can be applied generally in making other drawings. What an artist discovers in observing and drawing a particular object can lead to the development of a pictorial technique. As a technique for rendering a certain aspect of the visible world, it will be of use in creating drawings both from life and from the imagination. Its use will demonstrate knowledge of how things look in general, when it comes to that aspect of visual appearance. The pictures that employ the technique will impart knowledge of this aspect of visual appearance, by displaying it as an element of a visible scene.

Bibliography

Abell, Catharine.
- "Canny Resemblance." *Philosophical Review*, Vol. 118, No. 2 (2009). 183-223
- "Pictorial Realism." *Australasion Journal of Philosophy*, Vol. 85, No. 1 (2007). 1-17

Bantinaki, Katarina. "Pictorial Perception as Illusion." *British Journal of Aesthetics,* Vol. 47, No. 3 (2007). 268-297

Berger, John.
- "Painting and Time." In *The Sense of Sight*. Random House, 1985. (205-211)
- "The Basis of All Painting and Sculpture Is Drawing." In *Landscapes*. Verso, 2016. (27-32)

Black, Max.
- "Making Something Happen." In *Models and Metaphors: Studies in Language and Philosophy*. Cornell University Press, 1962.
- "Metaphor." In *Models and Metaphors: Studies in Language and Philosophy*. Cornell University Press, 1962.

Budd, Malcolm.
- "On Looking at a Picture." In *Aesthetic Essays*. Oxford University Press, 2009. (185-215)
- "The Look of a Picture." In *Aesthetic Essays*. Oxford University Press, 2009. (216-238)

Clarke, Thompson. "Seeing Surfaces and Seeing Physical Objects." In Black, Max (ed.), *Philosophy in America*. Allen & Unwin, 1965. pp. 98-114

Danto, Arthur. "Seeing and Showing." *The Journal of Aesthetics and Art Criticism*, Vol. 59, No. 1 (2001). 1-9

Davidson, Donald. "What Metaphors Mean." *Critical Inquiry*, Vol. 5, No. 1 (1978). (31-47)

Forbes, Graeme. *Attitude Problems: An Essay on Linguistic Intensionality*. Oxford University Press, 2006.

Freeland, Cynthia. *Portraits and Persons*. Oxford University Press, 2010.

Gombrich, E.H. *Art and Illusion: A Study in the Psychology of Pictorial Representation*. Princeton University Press, 1960.

Goodman, Nelson.
- *Languages of Art*. Hackett, 1976.
- *Ways of Worldmaking*. Hackett, 1978.

Grzankowski, Alex. "Pictures Have Propositional Content." *Review of Philosophy and Psychology*, Vol. 6, No. 1 (2015). 151-163

Harrison, Andrew. "The Limits of Twofoldness: A Defence of the Concept of Pictorial Thought." In *Richard Wollheim on the Art of Painting: Art as Representation and Expression*. Ed. Gerwen, Rob. Cambridge University Press, 2001. (39-57)

Hopkins, Robert.
- *Picture, Image, and Experience*. Cambridge University Press, 1998.
- "What Makes Representational Painting Truly Visual?" *Aristotelian Society Supplementary Volume*, vol. 77, no. 1 (2003). pp. 157-158

Hyman, John.
- "The Causal Theory of Perception." *Philosophical Quarterly*, Vol. 42, No. 168 (1992). 277-296
- "Vision and Power." *The Journal of Philosophy*, Vol. 91, No. 5 (1994). 236-262
- "Pictorial Art and Visual Experience." *British Journal of Aesthetics*, Vol. 40, No. 1 (2000). 21-45
- "Realism and Relativism in the Theory of Art." *Proceedings of the Aristotelian Society*, Vol. 105 (2005). 25-53
- "Subjectivism in the Theory of Art." *The Monist*, Vol. 86, No. 4 (2003). 676-701
- *The Objective Eye*. University of Chicago Press, 2006.
- "Depiction." *Royal Institute of Philosophy Supplement*, Vol. 71 (2012). 121-150

Kelly, Sean.
- "What Do We See (When We Do?)" *Philosophical Topics*, Vol. 27, No. 2 (1999). 107-128
- "Seeing Things in Merleau-Ponty." In Taylor Carman (ed.), *The Cambridge Companion to Merleau-Ponty*. Cambridge University Press, 2004. pp. 74-110.

Kenny, Anthony. "Abilities, Faculties, Powers and Dispositions." In *The Metaphysics of Mind*. Oxford University Press, 1992.

Lopes, Dominic.
- *Understanding Pictures*. Oxford University Press, 1996.
- "Pictorial Realism." *Journal of Aesthetics and Art Criticism*. Vol. 53 (1995). 277-85

Martin, M.G.F. "What's in a Look?" In ed. Nanay, Bence. *Perceiving the World*. Oxford University Press. (160-225)

Matherne, Samantha. "Merleau-Ponty on Style as the Key to Perceptual Presence and Constancy." *Journal of the History of Philosophy*, Vol. 55, No. 4 (2017). pp. 693-727

Maynard, Patrick.
- *The Engine of Visualization: Thinking Through Photography*. Cornell University Press, 1997.

- *Drawing Distinctions: The Varieties of Graphic Representation*. Cornell University Press, 2005.

Merleau-Ponty, Maurice. *Phenomenology of Perception*. Trans. Landes, Donald. Routledge, 2014.

McDowell, John. *Mind and World*. Harvard University Press, 1994.

Noë, Alva.
- *Action in Perception*. MIT Press, 2004.
- "Reply to Campbell, Martin, and Kelly." *Philosophy and Phenomenological Research*. Vol. 76, No. 3 (2008). pp. 691-706
- *Varieties of Presence*. Harvard University Press, 2012.
- *Strange Tools: Art and Human Nature*. Hill and Wang, 2015.

Panofsky, Erwin. *Perspective as Symbolic Form*. Zone Books, 1991 (1927).

Peacocke, Christopher. "Depiction." *Philosophical Review*, Vol. 96, No. 3 (1987). 383-410

Phillips, Antonia. "Drawing From Life." In eds. Hopkins, Jim and Savile, Anthony. *Psychoanalysis, Mind and Art*. Blackwell, 1992.

Podro, Michael. "Literalism and Truthfulness in Painting." *The British Journal of Aesthetics*, Vol. 50, No. 4 (2010). 457-468

Russell, Bertrand. "Appearance and Reality." In *Problems of Philosophy*. Global Grey, 1912.

Ryle, Gilbert. *The Concept of Mind*. University of Chicago Press, 1949.

Ed. Schear, Joseph K. *Mind, Reason, and Being-in-the-World: The McDowell-Dreyfus Debate*. Routledge, 2013.

Schier, Flint. *Deeper into Pictures: An Essay on Pictorial Representation*. Cambridge University Press, 1986.

Sibley, F.N. "Seeking, Scrutinizing and Seeing." *Mind*, Vol. 64, No. 256 (1955). 455-478

Siewert, Charles. "Is the Appearance of Shape Protean?" *PSYCHE*, Vol. 12 No. 3 (2006). 1-16

Snowdon, Paul. "Perception, Vision, and Causation." In *Disjunctivism: Contemporary Readings*. MIT Press, 2009.

Strawson, P.F.
- "Causation in Perception." In *Freedom and Resentment and Other Essays*. Routledge, 2008. (73-93) (Originally published 1974)
- "Perception and Its Objects." In *Philosophical Writings*. Oxford University Press, 2011. (125-145) (Originally published 1979)

Walton, Kendall.
- "Pictures and Make-Believe." *The Philosophical Review*, Vol. 82, No. 3 (Jul., 1973). 283-319
- "Transparent Pictures: On the Nature of Photographic Realism." *Critical Inquiry*, Vol. 11, No. 2 (1984). 246-277
- *Mimesis as Make-Believe: On the Foundations of the Representational Arts*. Harvard University Press, 1990.

Warnock, G.J. *Berkeley*. University of Notre Dame Press, 1982.

Wartofsky, Marx. "Picturing and Representing." In *Perception and Pictorial Representation*. Eds. Nodine, Calvin F. and Fisher, Dennis F. Praeger, 1979. pp. 272-283

Wittgenstein, Ludwig. *Philosophical Investigations*. Wiley-Blackwell, 2009. (Originally published 1953)

Wollheim, Richard.
- "Art and Illusion." *British Journal of Aesthetics*, Vol. 3, No. 1 (1963). 15-37
- "On Drawing an Object." *Art and The Mind*. Harvard University Press, 1974. pp. 3-30
- *Art and Its Objects: With Six Supplementary Essays*. Cambridge University Press, 1980.
- *Painting as an Art*. (The 1984 A. W. Mellon Lectures in the Fine Arts.) Princeton University Press, 1987.
- "On Pictorial Representation." *The Journal of Aesthetics and Art Criticism*. Vol. 56, No. 3 (1998). 217-226
- "What Makes Representational Painting Truly Visual?" *Proceedings of the Aristotelian Society*, Supplementary Volumes, Vol. 77 (2003). 131- 147
- "Why Is Drawing Interesting?" *The British Journal of Aesthetics*, Vol. 45, No. 1 (2005). 1-10
- "In Defense of Seeing-in." In *Looking Into Pictures*. MIT Press, 2003. pp. 3-15