

Refuting Skepticism with Style*

Elijah Millgram

Department of Philosophy
University of Utah
Salt Lake City UT 84112
elijah.millgram@gmail.com

August 10, 2016

In memory of Hilary Putnam

You are a brain in a vat. A mad scientist has kidnapped you, has surgically removed your brain from your body, and is keeping it—that is, you, or what is left of you—alive in a vat of nutrient fluid. He has connected your nerve endings to a computer that stimulates them so as to create the illusion that you are sitting in what you now think are your current surroundings, reading a paper laying out the amusing hypothesis that you are a brain in a vat. This story is the standard contemporary presentation of Cartesian skepticism: the worry that things would look just the same even if you were

*I'm grateful to Lori Alward, Sarah Buss, Steve Downes, Anne Eaton, Paul Haanstad, Sherri Irvin, Ram Neta, Lex Newman, John O'Dea, Susanna Siegel, Paul Thagard, Lauren Tillingshast, Nick Zangwill, Rachel Zuckert and panelists Stacie Friend, Daniel Jacobson, Bruce Glymour and Jonathan Weinberg for comments on earlier drafts, and, for helpful discussion, to Lanier Anderson, Sarah Buss, Ben Callard, Pepe Chang, Alice Crary, Edwin Frank, David Hills, Amy Johnson, Tamar Laddy, Krista Lawlor, Adria Quiñones, Guy Rohrbaugh, Rachel Shuh, Dina Stein, Kayley Vernallis and audiences at Vanderbilt University, an American Philosophical Association Pacific Division meeting, the Inland Northwest Philosophy Conference, Monash University, the University of Otago, Weber State University, the Aesthetics Anarchy conference at Indiana University, a meeting of the American Society for Aesthetics and the Ludwig-Maximilians-Universität München. Work on this paper was supported by a fellowship from the Center for Advanced Study in the Behavioral Sciences; I am grateful for the financial support provided through the Center by the Andrew W. Mellon Foundation.

a brain in a vat, and so that you cannot know that you aren't.¹ Is the skeptical doubt well-founded? Could you, unawares, be a brain in a vat?

Only, or so I shall argue, if your aesthetic sensibilities are wanting. (Only if you have no taste.) The usual varieties of skepticism about the external world are oblivious to the way in which the illusion to which you are, according to the hypothesis, being subjected, must be a work of art, and so oblivious also to the ways in which questions of style arise regarding it.

In skepticism about the external world, the internal world, and in particular, the appearances—the way things seem to you—are held constant, and the external world is allowed to vary; a causal account is given of how the varied external circumstances (on the one hand, the world as we think it is, on the other hand, the vat) could produce the fixed internal appearances. The appearances are evidently a pivot of the argument, and I mean ultimately to exploit that fact to derive a moral about a philosophical tradition distinguished by the pride of place which it gives to them. I have found as well that the architecture of the argument I develop here requires special emphasis; before proceeding to the main argument of the paper, it will be helpful to work through a brief warmup argument, one which will make the form of the central argument more visible. Accordingly, I want to begin with a question that philosophers seem to have dropped about three decades ago: what is a mere appearance, or, as the jargon of the time had it, a sense-datum?

1

Analytic philosophers let go of the notion of a sense-datum—roughly, an appearance of something else, which appearance might, as it happened, be freestanding of the thing putatively appearing—because arguments had been generated purporting to show that there were no such things. Just for instance: On the view that gave rise to Cartesian skepticism, all one ever sees are appearances, and so they must be lying around just about everywhere. Ordinary-language philosophers, such as J. L. Austin, issued what amounted to a standing challenge to show them even a single sense-datum, and it turned out that this was much harder than had been anticipated. When a bird appears in your garden, you do not see the appearance: you see the bird, and to think otherwise is a kind of grammatical confusion. When the time comes to justify your claim to be seeing a robin, you do

¹Putnam, 1981, pp. 5ff; Descartes, 1641/1979, Meditation 1.

not adduce the appearance, but the bird's red breast.² The idea that you could identify such freestanding appearances became so effectively and unforgivingly tagged the "Myth of the Given" that, half a century later, one could find such prominent philosophers as John McDowell treating a philosophical position's commitment to the 'given' as a decisive refutation of that position.³

Now there was something these arguments got right, namely, that sensory appearances that are able to stand free of their putative objects are not to be found lying around in most people's gardens. The arguments for that conclusion cannot have been fully satisfactory, however, because they would have proved too much; for instance, that such appearances cannot be found in most people's living rooms. But that conclusion was false for most of the twentieth century, and is false now: they can (or, more carefully, approximations to them can). Think of the spectrum of technologies along which one finds, towards the low end, black and white photographs, then, a little nearer the center, color photographs, then vinyl sound recordings, then digital sound recordings, then, moving now towards the higher end of the spectrum, video, movies, 3D movies, and virtual reality. Sense-data, properly understood, are not ordinarily found in nature, but that is because they are a technological achievement. When I play John Hiatt's *Slow Turning*, I am presented with sensory appearances whose content is: John Hiatt singing and playing his guitar, with accompanying drums, tambourine, bass, vocalists, and so on. But there is, of course, no Hiatt, no guitar, and none of the rest of it; in my living room there are only the freestanding auditory appearances, and the device that is producing them. (The content of the sensory appearance is, importantly, not the device that produces them, unless, e.g., and reverting now to an obsolete technology, the record is scratched.) The mad scientist's vat—the deceiving sense-datum generator of the skeptical story—is the eventual successor to the slide projector, the stereo, the Blu-ray player, and the Playstation or Xbox: the five-sense, completely surrounding, high-resolution home entertainment center. What these devices produce are progressively higher-quality sense-data: more realistic sound, more realistic images, and, eventually, at the limit of the technological series, *completely* realistic images. And I want to allow up front that that's possible, reserving however the question of just what it means.

The mistake made by philosophers of the fifties and sixties, then, was

²Austin, 1962; there is, perhaps surprisingly, a similar move in Heidegger, 1996, pp. 25–28.

³Sellars, 1956; McDowell, 1996. As we shall see shortly, a dislike for the 'given' is not a very good reason to reject sense-data; sense-data, it will turn out, are not 'given' at all.

thinking of sense-data magically, as though they were infinitely thin objects that inexplicably sprang into being between ourselves and the items we were observing (or where, perhaps, there was in fact no object we were observing). But nothing springs inexplicably into being in this way.⁴ And since objects with the characteristics of sense-data are too exotic to be a product of nature, they must be manufactured if we are going to have them at all. Of course, no real object is infinitely thin, even in the metaphorical sense at hand, but as the techniques for manufacturing them become more sophisticated, they become, as it were, thinner and thinner; that is to say, the material substrate (the photographic print, the gigantic 1950s television cabinet) becomes progressively less obtrusive, until, eventually, at the limit, what one sees is *just* the image. The crudely magical thinking of mid-twentieth-century analytic philosophers made the actual sense-data they saw every day (clumsy as they were even by current standards) philosophically invisible to them.

Every decade has its epitomizing oxymoron: the fifties had the Cold War, and the sixties, the Free Store; the seventies I suppose had Pet Rocks, and the eighties had Artificial Intelligence; the first decade of the new millenium gave us Jobless Recoveries, and this decade has given us PermaTemps. The epitomizing oxymoron of the nineties was Virtual Reality, and so perhaps we are better prepared than philosophers of the fifties to see what we have in front of us. If we are to take seriously the possibility of being deceived by sense-data generated by a mad scientist, or an evil demon, we need to think seriously about what that possibility comes to, given what sense-data actually are (and could be). Now of course we *are* deceived by sense-data, all too often: TV news, for instance, typically misrepresents the events of the day rather badly. But this is not the kind of deception that Cartesian

⁴Here is another way to see what the problem was. Philosophers in this tradition used to talk about seeing the front surfaces of objects, as opposed to the objects themselves: as though the front surfaces were somehow detachable from the objects they were surfaces of. Now ask yourself: what would it take *actually* to peel the front surface off an object, so that you could see just the surface, and not the object? In the case of a potato or a cucumber, there is a real answer: a potato peeler.

Philosophical interest in sense-data was fostered by their alleged incorrigibility, which was supposed to be guaranteed by their thinness; the idea was that the less there is to be wrong about, the less you can get wrong. Now that we can see what thinness really looks like, we can also see that incorrigibility does not follow from it: one often enough mishears what one has playing on the stereo. That is to say that sense-data, as I am describing them, do not play the logical or epistemological role for which they were wanted, and of course a natural old-school response would be that what I am describing are not sense-data at all. But there is nothing which could occupy that role, and the items whose history I have been recounting are the only candidates we have to bear the name.

skepticism has in mind. Rather, the worry is that you might be watching television, or a sophisticated virtual reality successor to television, without realizing that you were: and it is in this that the deception is supposed to consist.

There is a long history of misplaced credulity in this area that should make us disinclined to let pass stipulations like: “The illusion would be so perfect that you wouldn’t be able to tell.” In ancient times, Zeuxis was said to have painted grapes so lifelike that birds pecked at them; this is, in retrospect, obviously a fabrication.⁵ When wax records first appeared, enthusiasts marvelled that one could not distinguish recorded from live sound—a pronouncement at which the contemporary audiophile can only himself marvel. RCA’s logo, which depicted a dog staring down the speaker of a primitive record player, was captioned “His Master’s Voice”; but, as anyone with household animals knows, dogs and cats by and large simply ignore recorded sound.⁶ Early CDs, which we now hear as scarcely listenable, were said to be indistinguishable from the concert hall. There is an important question in this vicinity: what accounts for the repeated and enthusiastic willingness to believe, contrary to the direct evidence of one’s senses, that image could pass for reality? But I will put this question to one side for a while in favor of taking a different tack. To assess the Cartesian skeptical worry, we need to think not only about the reproduction device—the stereo or flat-screen monitor or virtual reality headset or the futuristic mad scientist’s vat—but about the programming for it.

2

Music boxes, and perhaps other ancestors of the devices we are considering, did not cleanly separate the playback mechanism from the recording (although by the time we have gotten to player pianos, that separation is already in place). But stereo systems, Blu-ray players, movie projectors and Nintendo sets are useless if programming is not provided for them: CDs, a broadcast, an already recorded videotape, or a game cartridge. If we are to have sense-data, we must have programming content. Now, as we progress

⁵Gombrich, 1969, p. 206; Kris and Kurz, 1979, recount this story at p. 62, and survey many similar legends. The explanation they advance is focused on the cultural role of the artist, and does not cover cases like those below, in which the mythologized artist does not figure.

⁶Which is not to say that they’re unaware of it. Domestic animals also ignore most television programming, but ‘video catnip’ is the feline equivalent of pornography, and can rivet a cat’s attention to the screen.

along the just-mentioned technological spectrum, at whose high end the contemporary Cartesian skeptic expects to find the perfectly convincing illusion, we notice that content becomes more and more expensive to manufacture. Snapshots, to be sure, are cheap, but then the skeptic does not expect us to be fooled by a snapshot. High fidelity sound recording requires studios, sound technicians, mixers, acoustically flat speakers, and much else. Movies require sets, booms, lighting designers, hair stylists, caterers, scriptwriters, editing rooms and editors, directors, production managers, and on and on, into budgets that practically never go below tens of thousands of dollars, and that sometimes run into the hundreds of millions.⁷ And we can expect virtual reality, as its level of visual gloss increases, to be more expensive than movies: movies need accommodate only the point of view chosen by the director, whereas virtual reality requires computing images on the fly, to handle points of view selected by the user.⁸ As sense-data come closer to the ideal needed by the Cartesian skeptic, the costs of producing them go up. And this deep economic fact about sense-data has consequences for what those sense-data turn out to be.

Think about movies for a moment. Because movies are so expensive to make, they must have large audiences, and, because audiences are disinclined to see unfamiliar faces, big budget movies are made with name actors. Movies are tightly edited, covering days, weeks, or months of events in about ninety minutes. They are scripted for narrative closure. They have soundtracks. Because the paying public is not the only component of the budget, movies contain product placements. And because everyone involved in making a

⁷See, for instance, Bushkin, 2005, a promotional documentary which details the resources used in making a 45-second scene in one of the *Star Wars* films, and which estimates that over seventy thousand man hours went into the clip.

⁸This point is nicely made by Dennett, 1991, pp. 3–7. Dennett’s argument, to the effect that generating a convincing virtual environment is not computationally tractable, can be tightened up somewhat.

An NP-complete problem is, intuitively, a problem whose solution is very hard to find, but once found, easy to check. (The non-deterministic Turing machine whose polynomial-time computation puts the problem in NP usually works by guessing all the possible solutions, and then checking them in parallel. So the hard part of the NP-complete problem, for the deterministic Turing machine, is finding its way to the solution in the first place.) There are bound to be NP-complete problems in computational geometry: given such-and-such a configuration of physical objects and light sources, how will the scene look to an observer? (E.g., what will a spoon in a glass of water standing in front of a mirror look like?) If the solutions to these problems are easy to check for correctness—that is, if they can be checked in polynomial time—but hard to find, then anyone with a spoon and a glass of water (or whatever the right problem turns out to be) will be able to verify empirically that he is not a victim of Cartesian skepticism.

movie has a professional interest in publicity, all but the very lowliest are given credits. And there is much, much more of the same.

So suppose you are a brain in a vat. If there is to be a convincing illusion, there must be professionally produced programming to play on the vat's reproduction device. That content has to have been made by a large team of highly skilled professionals using specialized and expensive equipment. Skilled professionals and equipment do not come for free, and this will have consequences very similar to those the analogous fact has for films. Your illusory life in the vat will be tightly edited. It will come with a musical score. When your spouse walks into the room, you will recognize him or her as Daniel Auteuil or Natalie Portman or some other famous actor. Brand new commercial products—computers, cars, breakfast cereals—will be overly conspicuous inhabitants of your world. And if there is any doubt whatsoever remaining, all you need do is wait ninety minutes, and there will be narrative closure, followed by very lengthy credits. In short, it would be very, very difficult to be a brain in a vat, and not to notice.

3

Thus far the warmup argument, and I expect a response along the following lines: To appeal to the economics of movie making is just not to take the skeptical hypothesis seriously! The hypothesis is that the image is *realistic*—that it is made to deceive you, and that we are if necessary to imagine that no expense is spared. I am shortly going to give that complaint a run for its money, but before I do, I want to take up the question of just what it is to take a skeptical hypothesis seriously, and in doing so, to put on the table what I take to be the methodologically principled response to it, and the correct frame for the subsequent argument: *which* skeptical hypothesis? Because when you put yourself in the business of arguing against such hypotheses, you have to choose your skepticism.

One way of putting the skeptical claim is that the world as it appears to be isn't *real*. But, as Austin pointed out, terms like “real” are always to be understood via some implicit contrast, and there are always many such contrasts that could be applied to an item at hand: money can fail to be real by being toy (Monopoly) money, or by being forged, or by being the previous government's since-discontinued currency, or in indefinitely many other ways.⁹ There is no such thing as showing that something is real *tout*

⁹Austin, 1962, pp. 70f; the idea is anticipated in Lewis, 1956, pp. 11, 15f.

court, and so any argument purporting to show that the world we see is real will have to address itself to one of those contrasts, and ignore the others.

The fairly tightly demarcated skepticism I mean to be addressing obeys two constraints. First, it allows manipulation of its victim's sensorium, but does not allow tampering with the metaphorical inside of his mind. The mad scientist operating the skeptic's vat is allowed to produce what you see, hear, feel, and so on, but cannot directly change what you believe, directly control your volition, or directly manage your train of thought. I suppose that the standard reason for excluding from consideration forms of skepticism under which your thinking may be, in one way or another, deranged, is that, without imposing this restriction, one cannot expect to make much headway. (If you aren't thinking straight, what good could *any* argument do you?¹⁰) My own reason, to foreshadow once more the destination of our discussion, is that I think we can learn something interesting about what it is to think philosophically about appearances by looking at the restricted form of skepticism.

That means that I am following much recent writing in departing from the skepticism of Descartes himself, who included insanity and dreaming, both of which alter the internal world, and in particular, one's cognitive capacities, under the skeptical heading. It means that I am not going to take up a recent and very interesting form of skepticism, one that denies that we even have the appearances of which we take ourselves to be aware.¹¹ And it means that I am putting Potemkin Village (or *Truman Show*) skepticism to one side—that is, the rather different skeptical worry that the part of the world that you're experiencing is, unbeknownst to you, something like a stage set.¹²

¹⁰Under this heading, let's just allow our envatted victim time to examine his world, to think about it, and even to cultivate an eye for it; it won't suffice for the skepticism we're considering if you can be fooled *for a moment*. And apropos volition, while attention is not, to be sure, fully voluntary, if one were not as free as we generally are (however free that is) to direct one's attention to different bits of the apparent world, whatever evidence of its manufacture that world contained would be all too easy to conceal. The ability to explore one's world with a certain amount of freedom will support a premise of an argument to come later, about the extent or size of the world we experience.

¹¹Dennett, 1991, pp. 354ff, 361f, recounts a laboratory experience that suggests that most of our apparent sensorium is our own interpolation, and I have been told that the compression algorithms that generate MP3 files rely on a similar phenomenon in the auditory domain.

¹²There is another spectrum stretched between the skepticism we are going to tackle, and this other sense in which the world might be unreal, and we need to be alert to the tradeoffs that we make as we move along it. The jennicam—a too-popular web site that showed still images of a woman's bedroom—was inexpensive to set up, and it let reality take care

Second, I am requiring the skepticism I am engaging to motivate radical doubt by supplying an account, made out in terms of our pre-skeptical take on the world, as to how the appearances could be as they are, while that of which they purport to be appearances might nonetheless fail to be as the appearances have it. A natural reason for insisting on what we can call a *motivated skepticism* is that only such a skepticism is compelling; if one shrugs off the need for the account to make sense in terms of our pre-skeptical take on the world, one might as well not tell the story of the vat in the first place, but just satisfy oneself with: perhaps I'm entirely wrong, who knows how? Such a skepticism, one might think, is hardly worth our time.

My own reason is that motivated skepticism provides us with philosophical opportunities that the alternatives do not. Motivated skepticism turns on *tangible* (rather than merely logical) possibilities. Logical possibility is not a very well understood notion, and the applicable bit of it amounts just to giving a consistent description. But consistent descriptions can be extremely minimal: “maybe you're wrong” is a consistent description. There is just not very much you can do in the way of thinking concretely about this sort of possibility, and so not very much that you can learn from refuting—or failing to refute—a skepticism of this sort. On the other hand, if we are giving a description, made out in terms of our current understanding of the world, that shows how we could be wrong about it, then we are laying out a skeptical possibility whose workings we more or less understand. If we understand a possibility in this more tangible way, then when we think about it, we may be in a position to draw substantial and surprising conclusions.¹³

of the detailing, but left little room for skepticism: what one was seeing was an actual bedroom in almost real time. As Anne Eaton has reminded me, Antonioni's *Blowup* shows us how unplanned details may intrude into photographs, even when the photographic reproduction is a stylistically artificial object. Movies can be complicated intermediate cases as well, depending how much they rely on documentary footage. Generally, the more of a representation is deferred to an automatic method of reproduction, the more difficult it will be to run the arguments I am developing here, but also the less room there will be to say that one is seeing something that is not really there; and that is true, in a somewhat different way, even when what is being reproduced is an object constructed to be represented via an automatic method of reproduction.

¹³Notice that the restriction puts omnipotence on the part of a would-be deceiver out of bounds: to make the author of the skeptical illusion omnipotent is really just to say that the illusion is created we-know-not-how; it is to give up on producing a skepticism driven by a tangible possibility. Familiarity with talk of omnipotence may have made this problem less obvious to Descartes' contemporaries.

The restriction also, perhaps more surprisingly, legitimates a presumption of motivational commensurability between ourselves and the hypothesized artist of the vat illusion.

If we in any case have to choose our skepticism, we ought to be very clear what we are choosing it *for*. We have had on the order of three hundred and fifty years during which refuting skepticism has been a rite of passage for the ambitious would-be well-rounded philosopher (the main requirement being that the execution be different than last time). It is not as though there is a shortage of perfectly good refutations of this or that skeptical challenge. Adding one more to the pile, at this late date, would be tiresome, were there not some further philosophical agenda to be advanced.

4

Let us return to the objection raised against the warmup argument, and suppose, for the sake of the argument, that one might be a brain in a vat, facing a lengthy bit of no-expenses-spared realistic super-cinema: the very obvious marks of manufacturing—the credits, the spectacular special effects, the cutting from scene to scene, the musical score, the well-known actors—have all been removed, in an effort to produce a convincing illusion, and to make your surroundings look as real as possible. What will the upshot of such an effort be?

The first thing to notice here is that the programming supplied for the vat must be something very much like a work of art. I don't here intend the worshipful implications that the term 'art' often has; it might well be very *bad* art. I mean just that it will be an object manufactured to look a certain way. Consequently, it will inevitably have aesthetic properties in the way that works of art do, and in particular, it will have a style. And that, in art, is just what realism is.

Think of realism in cinema. Decade after decade, the attempt is made to show things as they 'really look', and invariably, the result is an identifiable look and feel—often an aesthetic accomplishment of real note—that, from

It is not just that the would-be deceiver in the story has to have motives of which we can make (enough) sense for the story to spell out a tangible possibility. Given that there's a rough parity of capability between ourselves and the illusion maker, and that a creature of more or less our abilities wouldn't know how to begin going about deceiving us unless it was motivationally close enough to us to understand us from the inside, we can recover motivational commensurability as a further corollary. (I'm grateful to David Hills for pressing me on this point.)

Discussions of skepticism about the external world typically take it for granted that a counterargument relying on *a posteriori* knowledge about the external world necessarily begs the question. So notice that precisely the contrary requirement is entailed by the constraints we are adopting.

the distance of a decade or so, is quite obviously not the way things really look, or ever looked. Remember Vittorio de Sica's *The Bicycle Thief*, where realism came to, among other things, squallor and the futility of trying to stand up for your rights against a predatory world; it now looks like late-forties realism (or rather, Italian Neorealism with a Marxist edge), but not like the way things really look. Remember *The Diary of a Mad Housewife* and Warhol's *Trash*: both attempts, using very different techniques, to show New Yorkers as they *really* were (that is, this being the sixties, at their most unattractive); these now look merely like sixties realism. Remember *The Long Goodbye*, a movie that looks like what it is, a seventies movie, and a charming reminder of the Kurt Vonnegut accent that realism had in that decade—but not, by any stretch of the imagination, a movie that looks the way things *really* look.

If the realism of the appearances generated by the skeptic's vat turns out to be a style, we should expect the presence of style in nature (though not necessarily in the artifactual segment of the world: PT Cruisers were cars meant as a stylistic allusion to a comic strip, and were real nonetheless) to be a dead giveaway. But we are not yet home free. For one thing, forties realism looked real to the forties; it would be a weak and uncomfortable refutation of skepticism if we could only recognize the deception in distant retrospect.¹⁴ For another, unless we understand why attempts to reproduce the real manage (at the most) realism, we will not be in a position to insist that the effect is unavoidable.

5

We can make a start on the problem by taking up attempts of this kind in a somewhat different medium. *Trompe l'oeil* is a variety of visual realism, and the inability to be taken in by it is a failure of aesthetic responsiveness.¹⁵ What that means is that being taken in by it is not simply something that is

¹⁴Perhaps this restriction would still leave room for an inductive end-run around skepticism: if one is in a position to notice that one's previous surroundings must have been wholesale illusion, one may be able to learn to see through the current surroundings also. The cinematic realism of today—mumblecore, in one of its variants, appropriated into such mainstream movies as *Please Give* and *Tiny Furniture*—is already recognizable as a distinctive style.

¹⁵This very perceptive observation is due to Anne Eaton, and is seconded by Ebert-Schifferer, 2002, p. 24. N.B.: it suffices for our argument if the claim I am advancing is true for the most part; if your favorite Italian church really does manage straightforward architectural illusion, that does not impugn the present point about how *trompe l'oeil* normally operates.

done to one; it requires one's cooperation, and if it is deception, it is closer to self-deception than being the victim of simple misinformation. The usual label for the phenomenon is *suspension of disbelief*, and it is pervasive in the reception of aesthetic objects.¹⁶ Art high and low—from genre novels to Greek tragedy—succeeds by inducing its consumers to suspend disbelief, that is, not by deceiving them but by eliciting a more or less generous acquiescence. Now note that the willingness to go along is normally a response to, among other things, style—to the way, for instance, that the conventions of the relevant genre are being followed.¹⁷ Since this is no less true when

¹⁶The phenomenon is puzzling, and explicating it is one of the most persistent open problems in the philosophy of art. (For a characteristic expression of the puzzlement, one which tries to do away with it by denying that there is anything to be explained, see Schaper, 1978.) In thinking about the problem, it is important not to be misled by the surface of the phrase into assuming that it is anything as simple as voluntary belief; Williams, 1973, is good account of why this approach would be a nonstarter. And it equally important not to suppose that merely by relabeling it one has solved the problem; many of the prominent attempts to address it amount, in my own view, to no more than this.

¹⁷Some very thoughtful and sensitive discussion has recently been devoted to the so-called “puzzle of imaginative resistance:” while reading a novel, one has no difficulty in imagining various counterfactuals, but one cannot (or anyway, it is much harder to) imagine (let's call them) counterfactuals. Gendler, 2000, p. 58, writes: “I have a much easier time following an author's invitation to imagine that the earth is flat than I do following her invitation to imagine that murder is right;” likewise, Moran, 1994, p. 97: “We can easily ‘accept’ the existence of ghosts as one of the fictional truths implied by the story, but it seems we cannot similarly ‘accept’ it as true in that fictional world that, say, the murder of one's guest... is not to be condemned.”

But while I like bits of the explanations (and more on that in a moment), there is, if you look around at *our* planet, nothing to explain, and the repeated conviction that there is such a puzzle should be regarded rather as I earlier suggested we take assertions that you just can't tell the artwork from the reality—that is, as themselves incredible, and so, themselves a puzzle in aesthetics or moral psychology. As Landy, 2012, pp. 30f, has noticed, every competent reader of novels and every consumer of movies accepts counterfactuals just as routinely as counterfactuals. (E.g., when you read a vampire novel, you accept that vampires are evil and that killing them is good; as Laurell Hamilton has noticed, if there really were vampires, we would regard them as an ethnic minority with a congenital disability, which is to say that the counterfactuals effortlessly adopted by the consumers of such novels are on a par with bigotry of the worst kind.) Whatever Gendler and Moran are explaining, it cannot be their nonexistent explanandum.

Moran, discussing *Macbeth*, notices that style does a lot to get imagination off the ground; Gendler notices that the lack of imaginative engagement she wants to explain is a matter of unwillingness rather than inability. Putting these points together gives us both the explanation and the explanandum: recalcitrance at (often but not always) stylistic failure. Sometimes we are too impatient to go along with badly rendered descriptions of the factual flavor, other times we are too impatient to go along with badly rendered descriptions of the evaluative flavor, but we go along with both, about equally easily, when

the style is realism than when it is not, the realistic style must be apparent, on-the-spot and first-time-through. (We are awfully good at identifying style, and we do it awfully fast, but that claim should not be overread: it may take time, and a good deal of experience with a style, to come to see just what the constraints—and, correspondingly, just what the degrees of freedom—that constitute it are.)

To leave intact not only appearances but one's intellectual abilities means leaving intact the mastery of aesthetic notions such as style. You can imagine someone brought up without the advantages of art history classes, exposure to genre fiction, and so on; but our own experience is that we have had these, or similar, advantages, and the skeptical hypothesis we are considering keeps this fact a constant. However, this raises the question of why we should think that the envatted can have learned to recognize a style when they (we) see one: mightn't the mad scientist have exposed them (us) to misleading examples of styles in what they thought were their art history classes? And if, as I am about to argue, the vat illusion is *inevitably* stylized, wouldn't that entail that the examples from which we learned the notion of style were, equally inevitably, not real examples of a style, but rather stylized representations of styles?

So it is worth noticing that typically we do master the notion of style by way of what are acknowledgedly stylized representations of styles and toy instances of genres. The kind of formal education that one receives in art history classes is supported by the holdings of teaching museums, slide shows that present, for the most part, uncomplicated members of the type on display, and the like. What we learn the notion of style from are *anyway* stylized representations of styles, so we should not worry too much that, on the skeptical hypothesis, that is all we would have inside the vat. The ability

we are inclined to be generous to a rendering. And the ground for generosity is, again often but not always, the successful deployment of a style.

Recent discussions of Cartesian skepticism are themselves corroborating evidence for this last claim. When the time comes to invoke a cinematic portrayal of the skeptical hypothesis, we are invariably given *The Matrix*—a movie riddled with inconsistencies having to do with the skeptical hypothesis itself, and overall a very weak piece of science fiction. There are many other films that take up the same topic, and do so with much more intelligent attention to making the skepticism coherent: *Existenz* and *The Thirteenth Floor*, to name just two with release dates very close to that of *The Matrix*. So why do philosophers, of all people, prefer the incoherent rendition of their problem to intellectually superior alternatives? And how is it that they can suspend their objections both to the theoretical content, and to the film's juvenile (superhero-worshipping) values? Quite evidently, because the film is so *stylish*: the designer trenchcoats, the *faux*-retro cell phones, the terminally cool black sunglasses and all the rest trump whatever imaginative resistance even PhD-bearing philosophers might have mustered.

to cotton onto styles and genres that forms the basis of our education in the consumption of art is the ability we are relying on to allow us to discern the style of the illusory world of the skeptic's vat.

6

Still, if the ability to recognize style is to carry the argument, and if we want to avoid the kind of magical thinking that I earlier complained of in the sense-datum theorists, then we owe an explanation for it. To assemble one, I will begin with ingredients from the well-known work of E. H. Gombrich and unpublished writing by Andrew Hsu.¹⁸

Call artifacts with a function (what is sometimes called by philosophers these days a 'proper function') *devices*. Complex devices, Hsu points out, are composed of and depend on subdevices that also have proper functions. Now, while Hsu thinks that this fact is a specifically metaphysical truth, it suffices for our purposes that it's just the way the world works: you won't succeed in building a very complex functional device if you don't use as components 'smaller' functional—and functioning—subdevices. Generally, these subdevices will be members of artifactual kinds: a car has a carburetor as a component, and the carburetor is a type of artifact characterized by its function.¹⁹ Again, Hsu thinks of this fact as metaphysically unavoidable, because he thinks that devices belong to a functional kind in virtue of a history of members of the kind performing their function (for instance, my car has a carburetor because carburetors installed in previous cars actually worked). But it again suffices for our purposes that this is just the way the world is: without a long history of debugging, devices generally don't work, and a history of debugging entails there having been many instances of the kind that were actually *gotten* to work.

I said earlier that programming content for the vat must be something like a work of art, and so we can reformulate my gloss on that in the vocabulary I've just adopted as: a device with the (proper) function of appearing a certain way. We are supposing that the style of the attempted illusion is realistic, and we can further suppose that the artwork is an ambitious attempt at realism: unambitious attempts would have no hope of fooling their

¹⁸Gombrich, 1969, Hsu, n.d. We have here an illustration of a claim advanced in Millgram, 2013: it is typical of and instructive about philosophy that arguments conducted in such different subfields of the discipline can be made mutually supporting.

¹⁹'Generally,' because there can be one-of-a-kind devices; think of some of the big-ticket items on NASA's shopping list, and their unique, made-to-order components. I'm not sure how Hsu would want to handle such cases.

consumer, and the appearances produced in the vat are, jointly, supposed to constitute the envatted consumer's entire lived world. Consequently, the device—again, I mean the imagery produced by the vat, not the vat itself—must be extremely complex, and consequently, it must be assembled out of subdevices which themselves have (proper) functions; these subdevices must for the most part be members of artifactual kinds, and these kinds will have histories.

And in fact this is what we find whenever we see even moderately ambitious attempts at realism in representation: in realistic painting, just for instance, a subdevice for producing the appearance of a spinning wheel;²⁰ in the more recent genre of digital animation, a subdevice for producing weathered-looking surfaces;²¹ and I vividly recollect a video installation at a *Documenta* some years back, which showed a special effects technician explaining the workings of one such subdevice used in traditionally filmed movies: in distance shots of bodies going over cliffs, through car crashes, and so on, realism is achieved by wrapping full-size model skeletons with meat. (The technician recounted how popular the new trick became once it

²⁰Gombrich, 1969, pp. 228f; he supplies an illustration from Velázquez, and his treatment surveys many more such subdevices.

Michael Fried has devoted his career to documenting the history and problems of one such device, which he calls “absorption,” and because it is both exemplary, and has received such a lengthy treatment, let me talk through it quickly. The problem it is meant to solve has to do with a dilemma posed by the depiction of a human figure in a realistic painting: either the figure notices you, the viewer, or it doesn't. If it does, then the painting becomes stagey (“theatrical,” in Fried's way of speaking), and that impugns the warrant that you are observing what is there *anyway*, as opposed to a display being put on for you. (Here I will put to one side the question of why so many realist styles are committed to offering that warrant; when they're not, the skeptical question we are addressing does not arise.) If the figure *doesn't* notice you, and you are both right there, on the two sides of the fourth wall, an explanation is required for the illusion to be kept convincing. There are various solutions canvassed by Fried: the subject could be blind, for instance. The choice of subject matter on which he concentrates is a subject who is so ‘absorbed’—preoccupied, or just zoned out—that he or she plausibly fails to notice the viewer.

Absorption is a good generic example of the sort of realist device I have in mind, but notice that it will not be an element of the science-fiction realism that is supposed to underwrite external-world skepticism. The problem of not being noticed arises when the depicted figure fails to interact with a viewer. But the artworks being invoked by the skeptic are to be executed in media—such as videogames—that *are* interactive; the interactive world where everyone is too “absorbed” to notice you would be “theatrical,” not the other way around.

²¹Dorsey and Hanrahan, 2000; see also *Cinefex*, a trade journal devoted entirely to the nuts and bolts of CG work in commercial film, which describes many subdevices in this family.

was introduced, and how little time it took before everyone in the industry was using it.)

One class of realist subdevices is worth special mention, because it is a reminder that not all realistic techniques are illusionist, and because I will have a use for it down the road. To a large extent, realism works as a style by playing to its audience's preconceptions. Now, some but not all of one's preconceptions are such that seeing them represented provokes the response: "Yes, that's how things really are!"²² When a representation that prompts this response becomes cultural common currency, let's call it an *icon of realism*; a realistic style is typically marked by the use of icons of realism. While it is natural to describe the characteristic response to an icon of realism as recognition (and I will refer to it below as the 'recognition response'), of course such self-characterizations should not be taken at face value.^{23,24}

These subdevices, of whatever kind, inevitably have histories. They have to be invented and to go through a period of debugging, and it takes their users time to learn how to integrate them into larger and more complex

²²Of course, work in styles other than realism can prompt this response—think, say, of Kafka's novels. In these cases, a distinction is usually made between aspects of the work (e.g., medium or allegory, and content), and the reaction is focused on one rather than another aspect.

²³Icons of realism are often self-reinforcing. Splatter films, for instance, often strike their viewers as realistic in virtue of special effects depicting gruesome carnage, and it's worth remembering that most members of such a film's audience will never have seen gruesome carnage (or anyway, not seen it up close). Rather, they accept the representation as realistic because depictions in it resemble depictions previously accepted as icons of realism. E.g., the way a person is shown having his brains blown out is accepted as realistic because it looks like other cinematic depictions of a person having his brains blown out. A consequence of this sort of self-reinforcement is that realism tends to be clichéd.

²⁴This point reminds us that there are other aspects, and consequently other explanations (which does not mean *competing* ones), of the realistic style. The attempt to produce an object that looks like the way things really are turns out, necessarily, to amount to producing an object that looks like the way one *thinks* things really are. The realistic aesthetic object reflects, not reality, but its maker's preconceptions about reality (or anyway, its maker's view of his audience's preconceptions), and these preconceptions lend their imprint to the particular realistic style. However, *this* element of realism is indeed visible only in retrospect, its invisibility assured by Moore's Paradox. (That is, '*p*, but I don't believe it' is not just something like a contradiction; it's not really thinkable. Whereas, of course, '*p*, but I didn't used to believe it' is as thinkable as you like.) That recognizing this sort of subdevice requires one's preconceptions to have changed partly explains why a realistic style in cinema is so much easier to identify after the fact—even for viewers with well-developed aesthetic sensibilities.

projects.²⁵ So when we find large, complex and integrated devices, we will find institutions or subcultures with *long* histories—histories long enough to encompass the histories of the many subdevices. (The premise that the style is realistic is doing work here: many of the fleeting schools and movements of the twentieth-century art world restricted themselves to technical repertoires that needed only trivially short histories, but it is no accident that these sprang up with a rejection of realism.) So we can expect technically ambitious works of art—such as the illusory world of the vat—to bear the imprint of a lengthy and rich tradition.

There are many ways that the history of a developed representational craft must mark its products. For instance, the tradition will depend on the recruitment and education of generation after generation of fabricators. It will involve ideology of some kind (though not necessarily the same ideology throughout the tradition), serving to foster the *esprit de corps* needed to keep the enterprise going. It will train initiates where to apply the fabrication techniques they're taught—orientation of this kind being necessary because the integration of subdevices into larger devices has to be debugged, too, which means that the elements of the repertoire will be fit together, again and again, into the same few configurations (albeit with many varia-

²⁵Here again, absorption is a good example, just because Fried has been documenting its development, and the ways it has adapted to address technical problems it has posed for its users.

However, let me register a complaint about Fried's discussion of the tail end of this history. He notices that the device of absorption has been assimilated into recent photography, and takes Jeff Wall's work to be exemplary in this regard (Fried, 2008). It is not, and Fried is misunderstanding Wall's work. Absorption is a device that addresses a problem that arises within realist styles—see note 20, above. But Wall's work, the best of it, anyway, operates by undermining the realist style that it only apparently adopts. Fried himself observes that Wall's images are digitally assembled from different photographs, in a way which precludes the finished product having actually been seen by anyone; Fried does not discuss the ways in which the images are engineered to produce the gradually dawning realization that, as the viewer might put it, this couldn't be *real*. But the upshot is that Wall is not producing realist artworks, but rather, unobviously nonrealist works, works whose *subject* is realism, and so his appropriation of absorption cannot be taken at face value.

Fried titles the book in which he discusses photographers like Wall *Why Photography Matters as Art as Never Before*. If all that is going on is the adoption of one realist device from the painterly tradition, why does Fried think it matters *that* much? After all, it is, rather, *just* as before. My guess is that Fried takes the incoherent juxtaposition of *I can't see you*, as implicated by an elaborately worked object whose *raison d'être* is to be hung in a museum, to *make it Art*. This strikes me as a holdover of deconstruction, whose practitioners thought that pragmatic contradictions in a text made it *literature*. But here we are concerned only with art in the more minimal sense required by the argument, and I will not pursue the questions having to do with capital-A Art here.

tions). In works of art, this problem/solution pair will manifest itself as the tendency to genre—one of the more blatant ways in which style forces itself on the attention of even the barely attentive.

More importantly, within such a tradition, the repertoire of subdevices it deploys will have been shaped by a process analogous to natural selection. Subdevices that work will be reused and refined; those that don't will be abandoned and if possible replaced. Now such subdevices will count as working, for the purposes of selection, both when they induce suspension of disbelief, and when they pass for real in some more straightforward sense.²⁶ So we should expect many suspension-inducing subdevices to end up in the repertoire. Consequently, whenever we have an accomplished realistic style, we will find the artifacts it produces to operate (in part) by inducing suspension of disbelief. To reiterate, suspension of disbelief entails that the

²⁶Why can't the selection process discriminate between these two modes of effectiveness? It is an observation that, over the long term, it never has, and here are a couple of candidate explanations.

First, cultures sufficiently large and enduring to support the repertoires of an ambitious realism will have undergone ideological revolutions, in which, characteristically, some of the techniques, forms, dictates, and so on of a previous stage of their histories are dropped. (So realistic art will have, for instance, subject matter it does not handle, or handle well, not just because there is nothing in the bag of tricks that will do, or do it well, but also because some of the subdevices in the former bag of tricks have been proscribed, and are being *avoided*.) These revolutions seem inevitably to overthrow purist programs. For a recent but typical example, take Lars von Trier's 'Dogme 95,' which seems to have lasted less than ten years.

This is not the place to try for a full-fledged sociological explanation of that fact, so in place of that here is, perhaps more appropriately, an aesthetic explanation (due to Dan Jacobson). The realistic impact of deleting an unrealistic element to which we have grown accustomed is only temporary; once the recognition response to the deletion has been extinguished, realistic artists no longer see the point of insisting on it. (Perhaps they go hunting for another familiar convention to delete.) This may help account for a phenomenon I remarked on near the outset of the paper, the praise that early CDs received for their realism: the absence of ambient noise was, temporarily, an icon of realism. This may also be part—but only a small part—of the explanation for what Wollheim named “the improvisatory character of naturalism” (1987, p. 73).

Second, notice that there is an asymmetry between the demands imposed by the argument on artists and on art's consumers. In order to sniff out a would-be deception, the consumer need only identify a relatively small handful of devices as having the function of inducing suspension of disbelief; in order successfully to perpetrate such a deception against an alerted audience, the artistic tradition must root out all or almost all nonillusionist devices from its repertoire. But now notice further that the invitation to suspension of disbelief will operate on the originating artist and his colleagues as well; what reason do we have to think that *all* of such a tradition's emotional complicity in its own fabrications can be overcome? But if it cannot, why think that that the artists themselves can make the sufficiently thoroughgoing distinction between illusionist and suspension-inducing devices?

presence of style is in some way recognized. In short, a realistic style can invite and achieve near-instant suspension of disbelief (and so its instances must be registered—albeit subliminally—as styled by their consumers) because many of the subdevices that make up the repertoire of a realistic style are the product of a history of selection for precisely that.²⁷

The history shows that there is no real alternative to this gradual accumulation of technique. Theory-driven shortcuts, such as impressionism and pointilism, which attempted to reproduce the way things really look by invoking generic accounts of what the appearances must be (that is to say, painterly analogs of sense-datum theory), have uniformly produced what are among the most exaggerated and easy-to-recognize styles in the history of art.²⁸

The cumulative effect of the process I have been incompletely describing will not necessarily be all and only what has been denoted and connoted by the essentially contested concept of style.²⁹ But what matters is that we have on hand an explanation of (some substantial and recognizable component of) style's being both inevitable and detectable in an ambitious and realistic work of art. The programming content for the Cartesian skeptic's vat, it has been allowed, is both. I suggested earlier that, like other styles, realism

²⁷The move here has a model in Terence Deacon's rebuttal of Chomsky (1997, ch. 4). Why are natural languages so learnable? We do not need to appeal to hardwired neural modules with astonishing powers once we remember that languages are themselves products of a selection process, one that filters their components for learnability.

²⁸Programs such as impressionism fail because there is no finite vocabulary of appearances; it follows that the fabricators of a world such as our own, into which we can introduce novelty, must constantly invent further modes of appearance. (People build telescopes, or microscopes, or start making polyester clothing, and suddenly the mad scientists have to figure out how they are going to make such things *look*; Pixar's history shows this to be an occasion for genuine and impressive creativity.) This means that the vat's software must be regularly upgraded, to accommodate these added features. And as Microsoft's experience demonstrates, a constant stream of new features means an endless supply of new bugs—even when you have a programming team as large and as well paid as Microsoft's. This suggests a further reason for being fairly confident that we are not brains in vats: our perceptual world isn't constantly crashing.

²⁹And for that reason it will not necessarily respect distinctions belonging to alternative attempts to firm up the concept. Just for instance, Wollheim has insisted that "it is a defect in traditional connoisseurship that the distinction between style and signature has not always been recognized," and he distinguishes as well between individual and general style (1987, pp. 26ff, 36). But on my use of the notion these are deeply continuous with each other. Or again, it is now a commonplace that Putnam's externalist insights about content have an analog for much of what is called style; it is easy to imagine a Putnam/Burge-style narrative in which, say, Jules Kirschenbaum's twin has scarcely any of our Kirschenbaum's highly allusive style, because there is nothing for him to allude to. But these are not the aspects or elements of style on which the present argument turns.

normally succeeds not by deceiving, but by soliciting suspension of disbelief. We can now see that this hypothesis does not require the impossible of consumers of artworks. But where disbelief can be suspended, a self-aware viewer can suspend that suspension in turn. In a way, Descartes was right after all: the problem is not that you might be the passive victim of illusion, but that, faced with a temptation to misjudgment, you might *acquiesce*.

7

Let's break and consider three objections. First, suppose we grant that the vat-image must be styled, and that we are actually quite good at recognizing styles. Still, if that capacity is not magical, is it reasonable to expect recognition of a style on the basis of a single instance of it?³⁰ Isn't it long exposure to a genre that equips one to recognize instances? So the objection to our argument is that if one spends one's entire life in the vat, one will have seen only one instance of this kind of super-cinema, and so one will be unequipped to recognize it for what it is.

We can motivate the objection with an idea that Wollheim attributes to Gombrich, I'm not entirely sure how fairly. An important aspect of a style is its state space: the range of options available to an artist within the style.³¹ You can't tell what the full range of possibilities is by examining just one possibility, and so you can't grasp a style by seeing just one instance of it.

A quick reply is that you can often see at a glance *that* an object is styled, even if you cannot yet say what the state space of that style is; and that is all our argument requires. A sonnet is visibly a highly styled object, even if you are unprepared to explain, on hearing your first sonnet, what it takes to be one. But a less hasty reply is that when the instance is extensive and varied, extent can do duty for the multiple instances.³² Take a book-length poem—such as Vikram Seth's *The Golden Gate* (1987), which is composed entirely of sonnets. Even if you have never seen a sonnet when you open it up, by the time you are finished reading, you will have a very good idea of what

³⁰I'm grateful to David Hills for pressing me on this. We don't want to overstate the objection; certainly, some styles and genres are undemanding in this regard. One usually needs to hear only one light bulb joke to be able to make up light bulb jokes of one's own.

³¹Wollheim's (not Gombrich's) word is 'repertoire,' but I am already using that word for the set of subdevices available to the artist. His discussion is at Wollheim, 1980, secs. 28–31.

³²On one standard metaphysical view (not my own), we *do* tell what the full range of possibilities is—what the range of possible worlds is—without having seen anything other than one of them (the actual world).

the elements of a sonnet are, and what you can do with them. (Although a delicate and important qualification has to be appended to that: poets are always trying to push the boundaries of such forms, and so in some very real sense the edges of the state space are genuinely undetermined.) The world whose authenticity I am proposing we test by looking for traces of style is experientially enormous, its extent being mostly a matter of the indefinitely many degrees of freedom it evidently supports: when you turn a pot over, or look behind a door, or in a trunk, or under a bridge, there's always something there: unlike even the most overprogrammed videogames, the world doesn't just unexpectedly *give out*.³³ And so if it has a consistent style, the fact that we only experience one such world shouldn't get in the way of determining that, and of figuring out what the style is.

Next: Once again, goes the argument we have just seen, the vat-illusion must be a work of art. As a work of art, it will have an identifiable style. The real world observably does not (with the exception of the parts of it that are manufactured objects) have a style. Therefore the world that we see is not a vat-illusion. Our second objection to that argument is that the presence of style does not allow us to discriminate a real from an illusory world, because the real world is styled, too.

This objection comes in two versions, one due to Oscar Wilde, who claimed that Nature imitates Art, archly insisting that sunsets were simply second-rate Turners, haystacks in the light simply second-rate Monets.

³³Once again, there are further anti-skeptical arguments in the neighborhood, and let's gesture at two of them now, because they can give us more of a sense of what the realistic style comes to, and how it is to be contrasted with the real. First, if omnipotence is out of bounds, then so is the infinitely articulated and elaborated illusion. Because we don't know what it would be (like) to control such an infinitely filligreed appearance, the inexhaustible density of detail is a mark of the real. (This description comes from Goodman, 2001, p. 161; the point itself is older, and can be found in Ruskin, 1906, vol. I, pp. 254–56.) Of course, it is important that the detail we see is precisely not of the sort that can be generated automatically (in the way, say, that fractal-based algorithms produce background detail in videogames).

Second, because an infinite elaboration of detail is one of the hallmarks of reality, selectivity on the part of the artist is inevitable. It is simply not possible to notice everything; most of what one sees must be allowed to slip into the background. And so the manufacturer of an aesthetic object will be in a position to control the representation or reproduction of only those elements that he notices—that are, for him, in the foreground. Even the most heroic attempt at realism will have to limit itself to capturing only some features or aspects of reality. (As Barthes noticed, certain forms of realism achieve the realistic effect by piling up detail in apparently indiscriminate fashion; but even these are no exception.) A shift in focus on the part of the viewer will highlight the contrast between the aspects of the aesthetic object that are under control, and those that are not; a great part of any style will be a matter of *which* range of elements it attempts to control.

“Things are because we see them, and what we see, and how we see it, depends on the Arts that have influenced us.”³⁴ When we see the unmanufactured world as it is, we see it through the lenses of style, because our perceptions are formed by art. (And we do not necessarily see it falsely; Wilde is suggesting that our secondary-quality understanding of colors ought to be extended to much else.) Works of art have a style, but the real world resembles them in this respect. So we cannot use the presence of a style to tell us whether we are seeing reality or virtual reality.

But there is an unmistakable difference between the style of a Turner and of a sunset from Indian Rock. The contrast is nicely elicited by a Woody Allen short story whose protagonist is ‘projected’ into *Madame Bovary*; his appearance in the novel is noticed by professors who muse that “the mark of a classic is that you can reread it a thousand times and always find something new.”³⁵ And so you can, but not *that way*. Allen’s story is funny because the two modes of importing an element into an artwork are so blatantly different that confusing them would be unthinkable; we needn’t worry, then, that anyone *will* confuse them.

The other variation on the objection is that the world does have a style (or several of them) of its own, and one which is not the contribution of our perceptual efforts—evident, perhaps, in a peacock’s feathers, or the repeated use of bilateral symmetry in living organisms, or even in the regularities we summarize as physical laws. (We don’t normally notice it, but solely for the reasons we think only *other* people have accents.) Here we need to remember, one, that if the world we live in is a work of art, it must be *pervasively* styled. We have alternative ways of explaining candidates for stylistic markers, found in one or another segment of our world (say, explanations having to do with natural selection, in the segment consisting of living organisms). If the world does appear pervasively styled, we can avoid the conclusion of our argument only by generating such alternative explanations for *all* such segments of the world. But if we do produce alternative explanations for each segment of it, eventually, insisting that what we have before us is not a style but something else will become ad hoc resistance to the obvious. At that point, the time will have come simply to acknowledge the presence of style, and to embark on an inference to the best explanation. My own judgment, however, is that we are nowhere close to that point now. Two, recall that the sort of style on which the argument

³⁴Wilde, 1889/1986, pp. 78–80. This should not be confused with a distinct point, that our theoretical representations of the world, and in particular strong scientific work, exhibit a style, often personal. On this point, see Ravetz, 1979, p. 105n.

³⁵Allen, 1981.

turns involves a repertoire of devices that invite an audience to suspend disbelief. If *these* appear in even fairly restricted segments of the natural world, the world as a whole will have betrayed itself as styled.

Our third objection also comes in more than one variation. One is prompted by a fact emphasized by Wollheim: that the proper appreciation of a painting requires awareness of its medium. (If you are doing it right, your attention oscillates between the objects depicted by the painting, and such features as its brushwork.) But the vat is not supposed to foist that sort of awareness of a medium upon us: it is not as though our lived world were the afterlife of *What Dreams May Come* (an iffy movie that showed Robin Williams splashing around in the daubs of paint that made up the scenery). Alternatively, only artifacts are styled, in the relevant sense; but the argument we are considering requires us to determine that the vat-image is an artifact by determining that it is styled. To identify an object as styled, don't we need to know that it is an artifact, first? Alternatively again, to consider an object as a work of art normally requires being cued: we are prepared to pay attention to the film's stylistic features because we are sitting in a dark theater; we treat the painting as a work of art because it is hanging in a museum, in a frame. But the vat is not supposed to be providing us with cues of this sort.

But, first, Wollheim was making a point about the proper appreciation of a painting. We need not generalize the features of that enterprise to other sorts of artwork, and in any case, the argument does not demand that sort of appreciation from the vat-image's consumer: we were not asking the envatted victim to *savor* the image to the fullest. Second, while it is true that the question of whether the surrounding world is a work of art does have to be raised in order for questions of style to be considered, and while it is true that the cues which normally raise the question are not present, this very argument does raise the question: if you are now inside the vat, and you have gotten this far into the paper you are now reading, the suggestion that your lived world ought to be examined for stylistic markings has already been made to you. Finally for now, even if identifying a style entails awareness of an artifact whose style it is, and a medium in which the style appears, that connection need not amount to temporally or evidentially prior awareness. Compare: That something eats, reproduces and so on entails that it is a living creature.³⁶ But it does not follow that I must have identified something as a living creature before I can

³⁶On this point, see Thompson, 2008, pt. I; for a brief recap of his view, see Millgram, 2009.

tell if it is eating, reproducing, and so on. On the contrary, the normal way of telling that something is alive is to observe that it eats, reproduces, etc. How *else* would I decide that some unfamiliar being was a living creature? Why else would it even occur to me to ask?

8

Why *care* about skepticism? There are of course professional philosophers' reasons to care—traditionally, that the success of a skeptical argument tells us something about our standards for knowledge—and I will take up one of these in the final section of this paper. But the question is usually asked in the tone of voice of the wicked child in the Passover seder, who demands to know what difference it makes to *him*. After all, if the skeptical illusion is that good, what's missing?³⁷ We might frame the complaint as a disjunction: either the world is real, in which case believing in it is the right response, or it's not, in which case it's at least good enough to warrant suspending disbelief. Since what we go on to do will be the same either way, why does it matter which it is? The argument we have been developing puts us in a position to identify some of the generally overlooked costs of having skepticism come out true, and so I'd like to take a moment to explain what these are.

Sunsets can be breathtakingly beautiful in their pinks, oranges, and reds; a postcard of such a sunset is likely to be kitsch, and grating rather than beautiful. *Metropolitan* was a movie that realistically portrayed the social lives of upper-class East Side adolescents; by the near-perfection it brought to the depiction of uninspiring youths making uninteresting conversation, it occasioned critical enthusiasm that would have been completely out of place had it been directed to the film's subjects. And quite generally, the aesthetic responses appropriate to a representation or reproduction are often very different from those appropriate to its object or original. This means that if you are a brain in a vat, and you don't realize it, your aesthetic responses are bound to be systematically inappropriate.

The problem transforms itself, once you realize that you are living in a vat-produced image, rather than disappearing. We are able to see the vastness of the oceans or the cosmos as sublime because we think they really

³⁷Students are now routinely encouraged to take up this question in introductory philosophy classes, usually with Nozick, 1974, pp. 42–45, as the supporting text. Bouwsma, 1965, is the locus classicus of the argument that if it's good enough not to care, then the skepticism is incoherent.

are vast. By contrast, a while back, a remake of *Godzilla* was promoted by a misconceived advertising campaign that stressed the size of the title character. (Outdoor advertising carried such slogans as: “His foot is this big.”) *Godzilla*’s size, because it was merely stipulated, was unimpressive, and certainly not an occasion for the Kantian response.³⁸ If the world were an illusion, and we were aware of it, responses falling under the heading of the Kantian sublime would be systematically unavailable, and more generally, while our aesthetic responses might well be more appropriate to their circumstances, they would be for the most part ironic. There is nothing wrong with such responses on occasion, but a life lived solely in the postmodern pose is thereby a diminished life. The truth of the skeptical hypothesis comes with costs even after the skepticism is unmasked.³⁹

How high are these costs? That would depend on what the difference was between the response appropriate to the realistic style and the response appropriate to reality; so you might think that our assessment would have to be deferred until after we have determined what the elements of realism are; and you might not expect that to happen anytime soon. We have already seen how cinematic realism changes from decade to decade, and while the point that realism is a style has been made frequently, it is usually accompanied with characterizations of the style that differ sufficiently from one another to discourage the question of what the realistic style is and involves.⁴⁰ The problems here are genuine, and it is worth mentioning that

³⁸This may not, however, have been Kant’s own view of the matter. (I’m grateful to Rachel Zuckert here.)

³⁹There is an obvious way to try adapt the argument of this paper into a riposte to the argument from design: if God were the designer, the world would bear His signature style. Now I don’t wish to pursue the complications of a theological argument, and in particular I don’t want to worry about whether the density of detail that we have already mentioned might be taken for, precisely, God’s signature. But notice that the point we have just made bears on how we would respond to the world were we to realize it to be God’s artifact. For instance, the vastness of the cosmos might come to seem merely cutesy and overdone. One hears a good deal about the Problem of Evil, but not nearly enough about the Problem of Bad Taste.

⁴⁰For example, Barthes, 1989, highlights the use of detail that is superfluous from the point of view of the narrative, and Bryson, 1981, ch. 1, treats this technique as the opening move of a much more elaborate strategy in painting, which he identifies as realist. But in his discussion of the realistic (nineteenth-century) novel, Bersani, 1984, ch. 2, objects that this feature of Flaubert’s technique is nonstandard, and emphasizes instead “that secure sense of general design which allows us to relax over specific passages,” the “shared commitment to the portentous detail,” the way in which “time in realistic fiction is not merely chronological [but] shaped by a prior imagination of beginnings and ends,” and “the commitment to intelligible, ‘full’ characters, to historical versimilitude.” Or again, Auerbach, 2003, p. 491, picks out “the serious treatment of everyday reality, the rise

they ought to have philosophical consequences. For instance, the appeal to passages taken from realistic fiction has become a familiar feature of recent moral philosophy. But if realism is a style, and one whose effects and methods are as controversial as this, we should disallow such appeals—at any rate, until we are quite clear just what philosophical work is being shifted, in such arguments, onto the style, and whether the style is legitimately equipped to take it on. However, we can bypass this problem in at any rate an initial estimate of the damage a successful skepticism would inflict by turning to what I earlier called the ‘recognition response’.

If we were to try to characterize realism in terms of the response I suggested it characteristically invites or elicits, a review of icons of realism would naturally prompt the thought that ‘recognition’ is usually the expression of a posture of cynicism and disillusionment: that the realistic style is actually fantastic, that it is a fantasy of nastiness, and that it is gripping in virtue of what Nietzsche called *ressentiment*.⁴¹ I think trying to delineate realism by

of more extensive and socially inferior human groups to the position of subject matter for problematic-existential representation. . . the embedding of random persons and events in the general course of contemporary history [and] the fluid historical background” as “the foundations of modern realism”; his history of the emergence of realism in Western literature makes it out to be a matter of the ability of an elevated style to assimilate lowly subject matter; whereas, while Boyle, 1998, pp. 250f, 161 (attributing the point to Watt), agrees that “[r]ealism’ is a term that applies both to the way things are depicted and the way they are assessed,” but corrects Auerbach, observing that “the realistic novel. . . can represent anything except work.” Yet again, Diamond, 1991, pp. 40f, 51, describes realistic novels and stories as exhibiting features and facts rather than merely telling the reader that the fact or feature obtains; as fiction in which “certain things do not happen. . . pots do not talk, elves do not do chores while shoemakers sleep. . .”; and as driven by “a conception of how things work in our lives, what leads to what, what sorts of things do in actual fact determine how events proceed.” She contrasts realism with a hagiography in which we “are told that [Brother Pacificus saw the soul of his brother ascend direct to heaven], but we are not told at all what it was like to see such a thing.” However, Auerbach, 2003, p. 470, in a discussion of a passage from Balzac which is introduced as recognizably realistic, points out that the central feature of the situation described in it “is not established rationally but is presented as a striking and immediately apprehended state of things, purely suggestively, without any proof.”

Edwin Frank has suggested to me that, like “real” (in Austin’s treatment of the term; see note 9, above), “realism” is a ‘trousers word’: it relies on a contrast, often implicit, with some other particular style. (So to describe a work of art as realistic is to position it against, say, the baroque style in one instance, a classical style in another, and so on.)

⁴¹This approach to realism would not be unprecedented. Cavell, 1999, investigating skepticism, pressed a line of investigation very much in this methodological spirit. Emerson complained of materialism—the metaphysical realism of his day—that the desire to strip the world bare of its appearances was a form of voyeurism, and morally corrupt. (See, e.g., Emerson, 1960, p. 259.) And Nietzsche, who appropriated from Emerson the practice of couching epistemological claims in sexual metaphors, advanced a related accusation in

looking to this kind of emotional stake would be premature, as would using the emotional stake to explain the conviction produced by realistic works of art: stylistic success, whether the style is realistic or not, is the default and usually sufficient explanation for suspension of disbelief. But there is *something* to this thought, and to pursue it a bit further, I want to use a deservedly famous poem of Elizabeth Bishop's, "The Moose".⁴²

In the poem, realism appears as a conversation overheard on a long bus ride, about "what he said, what she said, who got pensioned; deaths, deaths and sicknesses," and Bishop goes on to show us the 'recognition response':

"Yes. . ." that peculiar
affirmative. "Yes. . ."
A sharp, indrawn breath,
half groan, half acceptance,
that means "Life's like that.
We know *it* (also death)."

She then allegorizes the intrusion of reality on realism by placing a moose in the path of the bus. The bus stops, passengers gawk and whisper, and the narrator asks:

Why, why do we feel
(we all feel) this sweet
sensation of joy?

Bishop has here got right the truth (and notice that I myself am now asking for the recognition response, but to a poem not executed in a realistic style) in our less restrained worries of a moment ago, about the bitterness and dourness that seem so often to accompany the indulgence in realism. She is exhibiting how different our characteristic responses to realism and to the real are, and how unfortunate it would be always to produce the one in place of the other.⁴³

The Gay Science (1887/1974, 2nd ed. Preface, sec. 4).

It is important to remember that we have not always insisted on realism's ugliness; in the seventeenth century, the tradition of Dutch still life did not identify the 'real' in 'realism' with the 'real' in *Realpolitik*. If one were to pursue this avenue of inquiry, one would want to address the question: why have we turned away from *that* sort of realism?

⁴²Bishop, 1983, pp. 169–173.

⁴³Style is not always noticed. Sometimes attention is drawn to it by beauty (or other aesthetic properties). (I'm grateful to David Hills for this point.) What this suggests is that a good way to make the deception we have been considering work would be to construct an illusory world that was *bad* art: ugly, downright boring, and so on. This means that we are epistemically fortunate (as well as fortunate in other ways) to live in a world as beautiful as this one.

9

What do we want to take away from yet another refutation of one of the most repeatedly refuted hypotheses in the history of philosophy? The punchline of this sort of external-world skepticism is not solipsism—it is not that there might *be* no external world, because what I was calling a motivated skepticism relies on our understanding of the ways the external world might vary—but rather, that we might be deeply wrong about what the external world is like. That is, its point is the radical underdetermination of theory by available evidence. Let me use the capitalized term, “Realism,” to mean, not now the style we have been discussing, but a familiar philosophical position.⁴⁴ Traditionally, Realism and Empiricism face off against each other, with their responses to skepticism about the external world serving as standard-bearers for their respective armies. The Realist insists that, no matter you will never be able to tell, there must be a fact of the matter—either there really *is* an external world, or it really *is* an illusion. The Empiricist responds that the Realist is talking rubbish: the point of one’s theory about the world is to ‘save the phenomena,’ and so the skeptical hypothesis is either strictly unintelligible or patently false. (The external world is a construction out of the phenomena, or a theory we ‘accept’ in order to predict the phenomena, . . .) The Realist, continues the Empiricist, is putting forth a distinction where there is no difference, and so presenting as a possible fact something that, one way or another, we do not have the resources to make out as one.

So the lesson I would like us to leave with has to do with Empiricism (Empiricism rather than Realism, because ‘Realism,’ used as a label for philosophical positions, is also a ‘trousers word’; the substance is all on the other side of the contrast).⁴⁵ Empiricism has been with us for a while, and looking over its history shows it to be a coin with two sides. The first is the commitment on which we just remarked, to saving the phenomena. The obverse of that commitment is a deep resistance—I mean that in something like the way it is meant in psychoanalysis—to thinking about the phenomena themselves. Hume took the phenomena to be impressions and ideas, the contents of which were given by a semantics of resemblance; even though his various skeptical arguments produced mutually inconsistent

⁴⁴Another direction to take would explore the connections between Realism, the philosophical position, and realism, the aesthetic style. For a beginning on the subject, see Diamond, 1991; however, by Realism in philosophy, she means something rather different from the view I have just mentioned.

⁴⁵See again note 9, and the last paragraph of note 40, above.

conclusions, he never turned back to reexamine the account of phenomena that was the shared premise of those arguments—and in this he epitomizes the Empiricist tradition.⁴⁶ Or to take a representative contemporary case, van Fraassen’s best-known work makes scientific theory responsible solely to the observations, and identifies the observations as “what the unaided eye discerns”. Here the distinction between what you can and can’t achieve with your body alone requires a philosophical account of the body, both to explain why this is a distinction that matters, and to explain what the distinction is in the first place. (Is a prosthetic eye part of your body? If not, why is the eye you were born with part of it? And if it is, why isn’t a telescope, or a microscope, also part of your body?) But no such account is on offer: once again, the Empiricist response is to take (what it construes as) the phenomena for granted, and resolutely not to think about them any further.⁴⁷

I have been trying to show what it would be like to look into Empiricism’s blind spot, and to think about the phenomena themselves rather than take them as brute. So what I have been doing is not meant in the first place as the defense of a philosophical thesis (although it is that), but rather as a philosophical demonstration. The Realist and the Empiricist share a willingness to take the appearances (the phenomena) as brute and unthought; they differ on how fastidious they are when it comes to what lies beyond the appearances. And so when they think about skepticism about the external world, they are all too quick to agree that we can stipulate that the appearances are kept fixed... and then go on to argue about how much sense we

⁴⁶Garrett, 1981, presents a typical instance of such an inconsistency.

⁴⁷van Fraassen, 1980, p. 59; see also pp. 16ff, 56ff, 214. In *The Mind-Body Problem* (1983), when Rebecca Goldstein’s character is asked what her dissertation is about, she uses the answer, “philosophy of body,” as a conversation-stopping joke. But it’s not a joke: another way to put the point is that when you make the distinction between what you can and can’t do with your body alone into the pivot of your philosophy of science, then you had better have some philosophy of body to go along with it. (Action theorists who argue about whether to identify ‘basic actions’ with bodily actions should take this moral to heart as well.)

In the spirit of the point I am making, Giere, 1999, p. 181 and note, points out a difficulty in substituting a *science* of body (physiology) for philosophy of body: a potentially vicious circle arising from the use of “psychology and physiology to tell us what are the observable substructures of our models of psychological and physiological systems”. van Fraassen, 2002, attempts to respond to the many complaints attracted by the earlier notion of observation, but in does so in a way that amounts to a further confirming instance of the claim I have just made about Empiricism; for supporting discussion, see Millgram, 2006.

That indictment of a grand tradition may sound like a stone cast by someone who thinks he is without sin. On the contrary, if I seem oversensitive about it, it’s because I’ve committed this very sin myself.

can make of the world outside radically varying, and how much we can know about it. When we think about the appearances, we see that we are not in a position to stipulate any such thing, and that, instead of a sterile debate about the reality of the external world, the skeptical hypothesis opens out into a discussion of, variously, style, and what is involved in producing a freestanding appearance. What we are going to get by leaving Empiricism behind is not Realism, but more interesting and instructive philosophy.

References

- Allen, W., 1981. The Kugelmass episode. In *Side Effects*, pages 59–78, Ballantine Books, New York.
- Auerbach, E., 2003. *Mimesis*. Princeton University Press, Princeton, 2nd edition. Translated by Willard R. Trask.
- Austin, J. L., 1962. *Sense and Sensibilia*. Oxford University Press, Oxford. Edited by G. J. Warnock.
- Barthes, R., 1989. The reality effect. In *The Rustle of Language*, pages 141–148, University of California Press, Berkeley. Translated by Richard Howard.
- Bersani, L., 1984. *A Future for Astyanax*. Columbia University Press, New York.
- Bishop, E., 1983. *Complete Poems 1927–1979*. Farrar Straus Giroux, New York.
- Bouwmsma, O. K., 1965. Descartes' evil genius. In *Philosophical Essays*, pages 85–97, University of Nebraska Press, Lincoln.
- Boyle, N., 1998. *Who Are We Now?* University of Notre Dame Press, Notre Dame.
- Bryson, N., 1981. *Word and Image: French Painting of the Ancien Régime*. Cambridge University Press, Cambridge.
- Bushkin, T., 2005. *Within a Minute*. Lucasfilm, San Rafael. Film. Released on the DVD edition of *Star Wars III: Revenge of the Sith*.
- Cavell, S., 1999. *The Claim of Reason*. Oxford University Press, Oxford.

- Deacon, T., 1997. *The Symbolic Species*. W. W. Norton, New York.
- Dennett, D., 1991. *Consciousness Explained*. Little, Brown, Boston.
- Descartes, R., 1641/1979. *Meditations on First Philosophy*. Hackett, Indianapolis. Translated by Donald Cress.
- Diamond, C., 1991. Realism and the realistic spirit. In *The Realistic Spirit*, pages 39–72, MIT Press, Cambridge, Mass.
- Dorsey, J. and Hanrahan, P., 2000. Digital materials and virtual weathering. *Scientific American*, 282(2), 64–71.
- Ebert-Schifferer, S., 2002. Trompe l’oeil: The underestimated trick. In *Deceptions and Illusions: Five Centuries of Trompe L’Oeil Painting*, pages 17–37, National Gallery of Art, Washington, DC.
- Emerson, R. W., 1960. *Selections*. Houghton Mifflin, Boston. Edited by Stephen Whicher.
- Fried, M., 2008. *Why Photography Matters as Art as Never Before*. Yale University Press, New Haven.
- Garrett, D., 1981. Hume’s self-doubts about personal identity. *Philosophical Review*, 90(3), 337–358.
- Gendler, T. S., 2000. The puzzle of imaginative resistance. *Journal of Philosophy*, 97(2), 55–81.
- Giere, R., 1999. *Science without Laws*. Chicago University Press, Chicago.
- Goldstein, R., 1983. *The Mind-Body Problem*. Laurel/Dell, London.
- Gombrich, E. H., 1969. *Art and Illusion*. Princeton University Press, Princeton. Second edition, revised.
- Goodman, L., 2001. *In Defense of Truth*. Humanity Books, Amherst.
- Heidegger, M., 1996. *Being and Time*. SUNY Press, Albany. Translated by Joan Stambaugh.
- Hsu, A. n.d. On the natural history of artifacts. Unpublished manuscript.
- Kris, E. and Kurz, O., 1979. *Legend, Myth, and Magic in the Image of the Artist*. Yale University Press, New Haven. Translated by Alastair Laing and Lottie Newman.

- Landy, J., 2012. *How to Do Things with Fictions*. Oxford University Press, Oxford.
- Lewis, C. I., 1956. *Mind and the World Order*. Dover, New York.
- McDowell, J., 1996. *Mind and World*. Harvard University Press, Cambridge, Mass.
- Millgram, E., 2006. Review of Bas van Fraassen, *The Empirical Stance*. *Philosophical Review*, 115(3), 404–408.
- Millgram, E., 2009. Review of Michael Thompson, *Life and Action*. *Analysis Reviews*, 69(3), 1–7.
- Millgram, E., 2013. Relativism, coherence, and the problems of philosophy. In O'Rourke, F., editor, *What Happened in and to Moral Philosophy in the Twentieth Century?*, pages 392–422, Notre Dame University Press, Notre Dame.
- Moran, R., 1994. The expression of feeling in imagination. *Philosophical Review*, 103(1), 75–106.
- Nietzsche, F., 1887/1974. *The Gay Science*. Random House, New York. Translated by Walter Kaufmann.
- Nozick, R., 1974. *Anarchy, State, and Utopia*. Basic Books, New York.
- Putnam, H., 1981. *Reason, Truth and History*. Cambridge University Press, Cambridge.
- Ravetz, J., 1979. *Scientific Knowledge and Its Social Problems*. Oxford University Press, New York.
- Ruskin, J., 1906. *Modern Painters*. George Allen, London.
- Schaper, E., 1978. Fiction and the suspension of disbelief. *British Journal of Aesthetics*, 18(1), 31–44.
- Sellars, W., 1956. Empiricism and the philosophy of mind. In Feigl, H. and Scriven, M., editors, *Minnesota Studies in the Philosophy of Science, Vol. I: The Foundations of Science and the Concepts of Psychology and Psychoanalysis*, pages 253–329, University of Minnesota Press, Minneapolis.
- Seth, V., 1987. *The Golden Gate*. Vintage, New York.

Thompson, M., 2008. *Life and Action*. Harvard University Press, Cambridge.

van Fraassen, B., 1980. *The Scientific Image*. Clarendon Press, Oxford.

van Fraassen, B., 2002. *The Empirical Stance*. Yale University Press, New Haven.

Wilde, O., 1889/1986. The decay of lying. In *De Profundis and Other Writings*, pages 55–87, Penguin, New York.

Williams, B., 1973. Deciding to believe. In *Problems of the Self: Philosophical Papers 1956-1972*, Cambridge University Press, Cambridge.

Wollheim, R., 1980. *Art and Its Objects (second edition)*. Cambridge University Press, Cambridge.

Wollheim, R., 1987. *Painting as an Art*. Princeton University Press, Princeton.