Objectivity not salience bias?
Commentary on Fischer, Engelhardt, Horvath, and Ohtani (2019)[1]

Justin Sytsma (VUW)

Eugen Fischer and colleagues expand on a body of empirical work offering a debunking explanation of a key assumption involved in the argument from illusion. Following Snowden (1992), we can distinguish between the *base case* and the *spreading step* in the argument. Fischer et al. target the base case. In the most prominent current versions of the argument, the key move in the base case involves the *phenomenal principle* (Robinson, 1994, 32): "If there sensibly appears to a subject to be something which possesses a particular sensible quality then there is something of which the subject is aware which does possess that sensible quality." In brief, Fischer et al. contend that the move here from a seemingly uncontroversial claim such as "the coin appears elliptical to me" to there being something of which the subject is aware that is elliptical requires that the initial claim be given a "literal interpretation" such that something elliptical has appeared to the subject. But they contend that under such an interpretation the claim should no longer be taken to be uncontroversial, assuming too much of what the argument needs to establish. And they argue that much of the intuitive appeal of this move can be explained in terms of accepting the claim based on the dominant usage of appearance verbs (e.g., I *think* the coin is elliptical), then shifting to the less salient phenomenal usage required for the conclusion. Fischer et al. then present the results of a series of nifty new studies in cross-cultural psycholinguistics to support the conclusion that people make stereotypical inferences warranted by the dominant sense of appearance verbs, even in contexts where this dominant sense is inappropriate.

As I said in a review of an earlier piece from Fischer, Engelhardt, and Herbelot (2015) on this topic, they "provide an excellent example of experimental philosophy in action" (2016, 304). And the same applies to the present article. In my opinion, this is a shining example of the promise of experimental philosophy, illustrating the sophistication and diversity of methods found in x-phi in recent years (see Fischer and Curtis 2019 for some examples). It also illustrates, I believe, the difficulty of doing experimental philosophy.

While I agree with much of Fischer et al.'s discussion, including their general concerns with the base case in the argument from illusion and the general plausibility of the Salience Bias Hypothesis (SBH) they call on (Fischer and Engelhardt 2019, in press), and while I am also an advocate of both critical ordinary language philosophy and evidential experimental philosophy[2], I am less convinced by their diagnosis of the appeal of the base case in the argument from illusion and their explanation of their empirical findings. I have two primary worries. First, I have doubts about the move from data on the quick judgments of lay people to a conclusion about philosophers' judgments. Second, I worry that there are simpler explanations of Fischer et al.'s findings.

1. Lay Judgments and Philosophical Judgments

Fischer et al.'s diagnosis of the appeal of the base case in the argument from illusion to philosophers strikes me as one plausible explanation, but I would have initially offered a different diagnosis. My explanation would have been that philosophers who feel the pull of the argument from illusion tend to hold a prior theoretical perspective on which sensory qualities are

[1] Thanks to Eugen Fischer for extremely helpful comments on an earlier draft of this commentary.
[2] Although I think the type of psycholinguistic methods Fischer et al. employ might also be fruitfully called on in "neutral" experimental philosophy (Sytsma and Livengood, 2016, 64-77).

associated with the mental states of the perceivers. This corresponds with an alternative possibility that Fischer et al. note:

> Arguably, the inferences at the root of arguments from illusion (and hallucination) only strike their proponents as so intuitively plausible because they believe from the outset in the existence of a complementary perceptual space, 'the mind', in which objects of awareness can be placed when evicted from the viewer's physical environment…. How this conception and contextually inappropriate stereotypical inferences interact to generate these paradoxes and the 'problem of perception' remains to be examined. (34)

I've found that many philosophers hold such a view.[3] Of course, this might owe in part to the influence of the arguments from illusion and hallucination. And both diagnoses could play a role, perhaps each driving intuitive plausibility judgments in different philosophers, or perhaps interacting in the development of philosophical views. Determining the origins of the judgments of philosophers about parts of a philosophical argument is a rather difficult task to say the least. And I don't think that Fischer et al.'s research (yet) gets at this, or that it is intended to get at this. Rather, I take them to be attempting to make the case that the SBH is a *plausible explanation* of the appeal of the inference at the root of the base case in the argument from illusion. To do so they present new empirical evidence of stereotype intrusion for appearance verbs in lay people.

How plausible one finds the SBH as a primary explanation of what is going on with philosophers is likely to depend on some details of Fischer et al.'s studies. The first reason for concern is that participants were instructed to "respond as quickly as you can" with Fischer et al. noting in the article that this was "to ensure responses in under 5s, before controlled processes could modify automatic cognition" (17). But one might expect that philosophers will tend to spend far longer than this thinking about the base case of the argument from illusion. This would not necessarily mean that stereotypical inferences do not play a role in the genesis of philosophers' judgments about the argument from illusion, but I think it does raise doubts about how readily we can move from findings on non-philosophers to a diagnosis of philosophers. A second reason for concern is that even when instructing lay people to respond as quickly as they could, the preference for the "is" texts in the visual pairs in the studies was not overly strong. It is such responses that Fischer et al. take to show problematic stereotype intrusion. Averaging across the visual items in their English-language study, however, participants only selected the "is" texts about 64% of the time. This is a relatively slight preference. If the first reason for concern is accurate—if philosophers do in fact think about the base case of the argument from illusion more carefully and over a greater period of time—then we would expect this preference to be smaller still. But then it is unclear that this would be sufficient to explain the philosophical appeal of the argument from illusion.

2. Issues with the Explanation

Focusing on just the evidence of stereotype intrusion in lay people, while Fischer et al.'s findings are consistent with their predictions, I see some difficulties for their interpretation and believe that there are simpler explanations that might be given. In laying out these difficulties, I'll focus on their first experiment, since I'm not competent with either German or Japanese.

Fischer et al.'s two main findings are that participants showed a consistent preference for the "is" texts in pairs with visual objects and that this preference was attenuated in pairs with

[3] And, interestingly, that lay people do not tend to hold such a view (e.g., Sytsma and Machery 2010, Sytsma 2012, Reuter and Sytsma in press).

non-visual objects. The explanation that they offer for these findings is somewhat complicated, and I must admit that I find it quite difficult. (Likely this reflects that I have no training in linguistics and am quite out of my depth here!) Nonetheless, let me try to walk through the reasoning for their predictions focusing on the two pairs of items provided in the text (with thanks to Eugen for his help here).

Let's start with the pair with a visual object:

6a. The hill seemed quite steep. The rambler thought it was gentle.
6b. The hill was quite steep. The rambler thought it was gentle.

Fischer et al. arrive at their prediction about this pair—that people will tend to prefer **6b** over **6a**—as follows. People will tend to assign the patient-role for the appearance verb in **6a** to the agent of the second sentence, reading it as "the hill seemed quite steep to the rambler." They will then tend to follow the I-heuristic, inferring from **6a** that the rambler thought that the hill was quite steep. Finally, they will then feel a tension between the two sentences in **6a**—one implying that the rambler thinks that the hill is quite steep, the other stating that the rambler thinks it is gentle—and because of this will show a preference for **6b**, where readers just need to suppose that the rambler got it wrong and there is no additional conflict between implied and stated belief-attributions.

Now let's turn to the pair with a non-visual object:

19a. The plan looked good. Cole believed it was terrible.
19b. The plan was good. Cole believed it was terrible.

Fischer et al. make a different prediction about pairs like this—that people will show an attenuated preference for **19b** over **19a** by comparison to what they predicted for pairs like **6a**/6**b**. The initial reasoning for this prediction is similar to that for the previous pair. Again Fischer et al. hold that people will initially tend to assign the patient-role for the appearance verb in **19a** to the agent of the second sentence, reading it as "the plan looked good to Cole." And they hold that people will then tend to follow the I-heuristic, inferring from this sentence that Cole thought that the plan was good. But Fischer et al. hold that the sentences are in such clear contradiction—one indicating that Cole thinks that the plan is good, the other that Cole thinks it is terrible—that people will now tend to re-assign the patient role for the appearance verb in **19a** from Cole to the author (e.g., "the plan looked good to me"), and will, again with the I-heuristic, infer that the author thought that the plan was good. Since this expresses an opinion the author might have more easily expressed by saying "the plan was good," Fischer et al. suggest that people will then tend to follow the M-heuristic, taking the author's use of "looked" in **19a**, as contrasted with the simpler "was" in **19b**, to indicate doubt about whether the plan really was good. And since participants aren't given enough information to decide whether Cole was so competent as to be able to see past the appearances to recognize that the plan was in fact terrible or was so incompetent that he was unable to recognize that the plan was in fact good, they will have a weaker tendency to prefer **19b** over **19a** than **6b** over **6a**.

I have two worries here. My first worry is that I don't find the sentences in **6a** (read in terms of the rambler filling the patient-role for the appearance verb and following the I-heuristic) to be notably less contradictory than the sentences in **19a** (read in terms of Cole filling the patient-role and again following the I-heuristic). It strikes me that there is a clear tension between the rambler thinking that the hill is quite steep and the rambler thinking that it is gentle, and that this tension is comparable to that between Cole thinking that the plan is good and Cole thinking that it is terrible. And both tensions could be resolved in similar ways, including by assigning the patient-role for the appearance verb to the author rather than the agent of the second sentence.

Whether other English speakers tend to be like me in finding the visual and non-visual texts to be similarly contradictory is an empirical question. To begin to explore this, I ran an online study directed at participants in the US.[4, 5] Each participant received a variation on the appearance verb texts from Fischer et al.'s 36 critical items, with the first sentence in each text being changed to assign the patient-role to the agent noted in the second sentence and following the inference from the I-heuristic, using either "thought" or "believed" in line with the second sentence of the text. For instance, for the sample pairs the following texts were used:

6c. The rambler thought the hill was quite steep. The rambler thought the hill was gentle.

19c. Cole believed the plan was good. Cole believed the plan was terrible.

The 36 texts were given in random order along with an attention check. For each text, participants were told that the sentences were about the same person and asked to indicate whether they thought there was a contradiction between the two sentences on a 1-7 scale with 1 anchored with "No Contradiction" and 7 anchored with "Complete Contradiction." Responses were collected from 42 participants age 16 or older who passed the attention check and reported that English was their native language. In line with my worry, the mean for each of the 36 texts was in the upper half of the scale, ranging from a low of 4.21 ("Sam thought their efforts were idealistic. Sam thought their efforts were self-serving.") to a high of 6.64 ("The jury foreman believed the accused was guilty. The jury foreman believed the accused was innocent."). Further, the mean of the average ratings for both the visual (M=5.57, SD=1.19) and the non-visual texts (M=5.83, SD=1.01) were significantly above the mid-point, and the same held for the texts that originally used "appeared" (visual: M=5.79, SD=1.31; non-visual: M=5.80, SD=0.98), "seemed" (visual: M=5.44, SD=1.30; non-visual: M=5.73, SD=1.26), and "looked" (visual: M=5.47, SD=1.30; non-visual: M=5.97, SD=1.12).[6] In potential support of Fischer et al.'s explanation, however, I found that the mean for the non-visual texts was significantly higher than for the visual texts, although the difference was small.[7] Similarly, the mean for the non-visual texts that originally used "seemed" and that originally used "looked" were significantly higher than their visual counterparts, although again the differences were small.[8] No significant difference was found for the texts that originally used "appeared."[9] Overall, it is at best unclear whether the differences seen here are large enough to explain Fischer et al.'s original findings.[10]

My second worry concerns Fischer et al.'s assumption that participants will initially tend to assign the patient-role for the appearance verb to the agent of the second sentence and won't

[4] It should be noted that Fischer et al.'s study was run on students from the UK, and it is possible that these populations differ in their judgments. Given the tasks and that both samples were comprised of native English-speakers, however, we would not expect a notable difference here.

[5] Participants were recruited via advertising for a free personality test through Google Ads. Participants were given a 10-item Big Five personality inventory after completing the philosophical questions.

[6] Visual: $t(41)=8.5567$, $p=1.2e^{-10}$; Non-visual: $t(41)=11.753$, $p=1.0e^{-14}$; Appeared, Visual: $t(41)=8.8205$, $p=5.1e^{-11}$; Appeared, Non-visual: $t(41)=11.94$, $p=6.3e^{-15}$; Seemed, Visual: $t(41)=7.1764$, $p=9.3e^{-9}$; Seemed, Non-visual: $t(41)=8.8817$, $p=4.2e^{-11}$; Looked, Visual: $t(41)=7.2961$, $p=6.3e^{-9}$; Looked, Non-visual: $t(41)=11.362$, $p=3.0e^{-14}$

[7] $t(41)=-3.0683$, $p=0.0038$, Cohen's $d=-0.47$ (small)

[8] Seemed: $t(41)=-2.24$, $p=0.031$, Cohen's $d=-0.35$ (small); Looked: $t(41)=-3.1641$, $p=0.0029$, Cohen's $d=-0.49$ (small);

[9] Appeared: $t(41)=-0.053743$, $p=0.96$

[10] It is perhaps worth noting that there was also a significant difference between the ratings for the texts using "thought" (M=5.52, SD=1.12) and the texts using "believed" (M=5.90, SD=1.06): $t(41)=-5.5302$, $p=2.0e^{-6}$. There was no notable difference between the means for the visual texts using "thought" (M=5.54, SD=1.22) and "believed" (M=5.56, SD=1.23): $t(41)=-0.3065$, $p=0.76$. But there was a significant difference between the means for the non-visual texts using "thought" (M=5.49, SD=1.15) and "believed" (M=6.11, SD=0.98): $t(41)=-5.9099$, $p=5.8e^{-7}$.

tend to reassign it to the author for the texts with a visual object. My first reason for doubt is that the texts were presented in pairs and participants were instructed to answer the question "only once you have read **both** short texts (a *and* b) in each pair." Doing so, I would expect participants to tend to treat *both* texts as having the same basic context, including that participants would tend to read the first sentence of each text as stating the judgment of the same person. Since the verb in the first sentence of the "is" items lacks a patient-role, Fischer et al. expect that participants will take the author to be endorsing the assertion. But if this is correct, and if people tend to consider the texts together as a pair, it would seem to make the author a salient option for the patient-role of the appearance verb in the other text.

My second reason for doubt is that even focusing on just the appearance verb text of the visual pairs, it strikes me as quite plausible that people would tend to read the assessment in the first sentence as stating how things appeared to the author rather than the subject of the second of the second sentence. Consider a text like **6a**: the two-sentence structure with the subject specified in the second sentence but not the first suggests to me that the author intends to be putting forward their own assessment in the first sentence, then noting a dissenting opinion in the second. If the first sentence was meant to state the rambler's impression, I would expect the author to indicate this more clearly, noting the rambler in the first sentence and either mitigating or flagging the potential tension. For instance, the author might have written "although the hill seemed quite steep to the rambler, she nonetheless thought it was gentle" or "perhaps surprisingly, the hill seemed quite steep to the rambler even though she thought it was gentle." Given such possibilities, that the author used two sentences and didn't mention the rambler until the second sentence, suggests that she was not stating the rambler's assessment in the first sentence, but her own.

Of course, that people will be inclined to read the first sentence of texts like **6a** as expressing the assessment of the author is an empirical claim. To test this, I ran a second study using the same recruitment strategy as before. Each participant received just the appearance verb texts from Fischer et al.'s 36 critical items in random order. For each one they were asked to answer who they thought was making the claim in the first sentence, selecting either the person named in the second sentence or the author.[11] Responses were collected from 44 participants age 16 or older who reported that English was their native language. For each of the 36 texts, a large majority of participants selected "The Author" (the lowest proportion was 0.75 for "The cat darting across the street looked small. Jim believed it was large."; the highest percentage was 0.91 for "Their relationship appeared rocky. Susan believed it was smooth." and "The building seemed grand. The passer-by believed it was unimposing."). Overall, 82.8% of participants selected "The Author" for the visual texts, compared to 83.0% of participants for the non-visual texts. Following the analysis employed by Fischer et al., for each participant I calculated the proportion of "The Author" responses out of six for each of the six classes of texts, then calculated the mean of those proportions. Each mean was significantly greater than 0.5.[12, 13] Further, the means for the texts using "appeared" with a visual object and with a non-visual object were not significantly different, and similarly for "seemed" and "looked."[14]

---

[11] For instance, for text 6a participants were asked "Do you interpret the first sentence in this text in terms of the rambler finding that the hill seemed quite steep or in terms of the author finding that the hill seemed quite steep?" and answered by selecting either "The Rambler" or "The Author."

[12] Fischer et al. hold that by aggregating across the six questions for each type of pair, they're measuring consistency of preference and that this then allows for the use of parametric tests. I follow suit in my analysis here, but I have concerns about whether parametric tests are in fact warranted or whether a non-parametric test would be more appropriate here. I wish to thank Jonathan Livengood for discussion with regard to the statistical tests.

[13] Appeared, Visual: $M=0.81$, $SD=0.26$, $t(43)=8.1072$, $p=1.7e^{-10}$; Appeared, Non-visual: $M=0.86$, $SD=0.25$, $t(43)=9.3308$, $p=3.4e^{-12}$; Seemed, Visual: $M=0.84$, $SD=0.27$, $t(43)=8.3363$, $p=8.0e^{-11}$; Seemed, Non-visual: $M=0.82$, $SD=0.24$, $t(43)=8.9734$, $p=1.0e^{-11}$; Looked, Visual: $M=0.83$, $SD=0.28$, $t(43)=7.8194$, $p=4.3e^{-10}$; Looked, Non-visual: $M=0.81$, $SD=0.26$, $t(43)=8.0291$, $p=2.2e^{-10}$

[14] Appeared: $t(43)=-1.8581$, $p=0.070$; Seemed: $t(43)=0.90263$, $p=0.37$; Looked: $t(43)=0.4957$, $p=0.62$

The results of this study suggest against a key assumption in Fischer et al.'s explanation of their findings. As noted above, the reasoning behind their first prediction assumed that participants would draw a stereotypical inference from the first sentence of these texts to *S thinks that X is F*, where S is the person noted in the second sentence. But participants overwhelmingly indicated that they instead took the author to be making the assessment in the first sentence.

3. Alternative Explanations

If the above study is accurate—if people overwhelmingly tend to read the first sentence of the appearance verb texts as reporting the assessment of the author rather than the person noted in the second sentence—then what explains Fischer et al.'s pattern of results? Absent a compelling alternative explanation, we might suspect that something is going wrong in my study, perhaps inferring that the questions were misleading or that the context provided by the "is" text in the pairs (somehow) shifts participants toward reading the first sentence of the appearance verb texts in terms of the person named in the second sentence. Fortunately, I think it is possible to construct alternative explanations of Fischer et al.'s results that are consistent with the worries raised in the previous section. Specifically, it seems to me that there are a number of differences on average between visual objects and non-visual objects of the sorts that Fischer et al. used in their questionnaire, as well as differences in the way we typically talk about them. And such differences might explain the preference for the appearance verb text in visual pairs but not non-visual pairs.

One possible difference that strikes me in looking through the items is that whether or not the visual objects actually had the properties at issue is more objective (more a matter of fact) while whether or not the non-visual objects had the property at issue is more subjective (more a matter of opinion). For instance, it strikes me as somewhat more objective whether a hill is quite steep than whether a plan is good. This possibility fuels a rather simple alternative explanation of Fischer et al.'s pattern of results: participants tended to favor the "is" text when they thought that whether X is F is more objective and to favor the appearance verb item when they thought that whether X is F is more of an opinion. Assuming participants tend to assign the patient-role for the appearance verb to the author, the task would then seem to be to decide whether it is better for the author to make a weaker claim (e.g., she *thinks* that the hill is quite steep, she *thinks* that the plan is good) or a stronger claim (e.g., that the hill *is* quite steep, that the plan *is* good). We would then expect participants to show a greater preference for the stronger claim when the author is in a better position to claim with certainty that X is F. But one factor in this will be whether X's being F is typically taken to be an objective fact or a matter of opinion.[15]

To assess the initial plausibility of this explanation, I ran a third study using the same recruitment strategy as before, but this time I gave participants just the first sentence from the "is" items for each of Fischer et al.'s 36 critical items in addition to an attention check in random order, and asked them to "select the degree to which it expresses an objective fact or a subjective opinion on a scale of 1-7, where 1 indicates that it is completely objective and 7 indicates that it is completely subjective." Responses were collected from 43 participants using the same

[15] Of course, judgments about the objectivity of X being F is unlikely to be the only factor that matters for assessing how good of a position the author is in to claim that X is F. For instance, some objective facts are more difficult to discern from surface impressions than other. Further, the second sentence in the items in a pair might be seen as making the possibility of error more salient. To illustrate, while we might think that there is a fact of the matter about whether the icing on a cake is pink (**45**), and that looking at it is typically a good way to discern whether the icing is in fact pink, noting that somebody disagrees would raise the possibility that the lighting conditions are unusual, generating doubt about the initial assessment.

restrictions as in the first study. In line with my impression of the sentences, the mean rating varied notably across the texts, ranging from a low of 2.23 ("the bird was a Hammerkop") to a high of 5.53 ("the young artist was talented"). And the ratings varied between the visual and non-visual sentences, with the mean of the average rating for each individual for the visual sentences being significantly lower (M=3.25, SD=0.67) than for the non-visual sentences (M=4.40, SD=0.67).[16] Further, the mean of the average ratings for the visual sentences was significantly below the neutral point, while the mean of the average ratings for the non-visual sentences was significantly above the neutral point.[17] In other words, participants tended to treat the visual sentences on average as being more objective and the non-visual sentences on average as being more subjective. The same held when looking just at the "appeared" sentences, just at the "seemed" sentences, and just at the "looked" sentences.[18, 19] While these findings do not establish that the pattern of results found by Fischer et al. owes to differences in how likely participants were to think that X being F was objective, and while further testing is needed, the findings do offer initial reason to think that this might well be a plausible explanation of the results.

[16] $t(42)=-8.5395$, $p=1.0e^{-10}$

[17] Visual: $t(42)=-7.3164$, $p=5.1e^{-9}$; Non-visual: $t(42)=3.9255$, $p=0.00032$

[18] Appeared, Visual: M=3.07, SD=1.01; Appeared, Non-visual: M=4.45, SD=0.84; $t(42)=-6.0909$, $p=2.9e^{-7}$; Seemed, Visual: M=3.67, SD=0.79; Seemed, Non-visual: M=4.50, SD=0.89; $t(42)=-5.4217$, $p=2.7e^{-6}$; Looked, Visual: M=3.01, SD=0.85; Looked, Non-visual: M=4.26, SD=0.70; $t(42)=-9.0263$, $p=2.2e^{-11}$

[19] Appeared, Visual: $t(42)=-6.0926$, $p=2.9e^{-7}$; Appeared, Non-visual: $t(42)=3.4984$, $p=0.0011$; Seemed, Visual: $t(42)=-2.7299$, $p=0.0092$; Seemed, Non-visual: $t(42)=3.7012$, $p=0.00062$; Looked, Visual: t(42)-7.6218=, $p=1.9e^{-9}$; Looked, Non-visual: $t(42)=2.4431$, $p=0.019$

References

Fischer, Eugen and Mark Curtis (2019). *Methodological Advances in Experimental Philosophy*. London: Bloomsbury.

Fischer, Eugen and Paul Engelhardt (2019). "Eyes as windows to minds: Psycholinguistics for experimental philosophy," in E. Fischer and M. Curtis (Eds.), *Methodological Advances in Experimental Philosophy* (pp. 43-100). London: Bloomsbury.

Fischer, Eugen and Paul Engelhardt (in press). "Living Stereotypes: Salience Bias in Philosophical Arguments," *Mind and Language*.

Fischer, Eugen, Paul Engelhardt, and Aurelie Herbelot (2015). "Intuitions and Illusions: From Experiment and Explanation to Assessment," in E. Fischer and J. Collins (Eds.), *Experimental Philosophy, Rationalism and Naturalism* (pp. 259-292). London: Routledge.

Reuter, Kevin and Justin Sytsma (in press). "Unfelt Pain," *Synthese*.

Robinson, Howard (1994). *Perception*. London: Routledge.

Snowden, P. F. (1992). "How to Interpret 'Direct Perception'," in T. Crane (ed.), *The Contents of Experience* (pp. 48-78), Cambridge: Cambridge University Press.

Sytsma, Justin (2012). "Revisiting the Valence Account," *Philosophical Topics*, 40(2): 179-198.

Sytsma, Justin (2016). "Rethinking the Scope of Experimental Philosophy," *Metascience*, 25: 301-304.

Sytsma, Justin and Jonathan Livengood (2016). *The Theory and Practice of Experimental Philosophy*. Peterborough: Broadview.

Sytsma, Justin and Edouard Machery (2010). "Two Conceptions of Subjective Experience," *Philosophical Studies*, 151(2): 299-327.