



# Perceptual malleability: attention, imagination, and objectivity

A Reply to Commentators: Kind, O'Callaghan, and Wu

Dustin Stokes<sup>1</sup>

Accepted: 15 July 2023

© The Author(s), under exclusive licence to Springer Nature B.V. 2023

## Abstract

This article offers a reply to commentaries from Amy Kind, Casey O'Callaghan, and Wayne Wu. It features a defense and further analysis of perceptual malleability, as defended in *Thinking and Perceiving*. In turn, it considers the consequences of malleability for attention and the cognitive penetrability of perception, imagination and perceptual skills, and perceptual content and objectivity.

**Keywords** Malleability · Modularity · Cognitive influence on perception · Attention · Imagination · Perceptual content · Objectivity

I could not be more grateful for such inciteful and productive criticism from three researchers whose work I hold in the highest regard. With their different points of emphasis, my commentators manage to cover quite a lot of the book (and the book covers quite a lot of ground). They also manage to extend and complement some of my explanations and arguments, which encourages optimism about future discussion and work. Finally, they underscore a difficult balance, one that I attempt to strike in the book and that will become clear as I work through my replies.

Let's begin with a few old friends. In a debate with Jerry Fodor about theory-ladenness and observation, Paul Churchland argues that bistable figures such as the Duck-Rabbit provide evidence for the cognitive penetrability of perception. He suggests that when one changes one's "assumptions" about the figure, one's visual experience of that figure changes accordingly (Churchland 1988). Fodor replies,

"False. One doesn't get the duck rabbit (or the Necker Cube) to flip by 'changing one's assumptions'; one does it by (for example), changing one's fixation

---

✉ Dustin Stokes  
dustin.stokes@utah.edu

<sup>1</sup> Department of Philosophy, University of Utah, Salt Lake City, USA

point. Believing that it's a duck doesn't help you see it as one; wanting to see it as a duck doesn't help much either. But knowing where to fixate can help" (Fodor 1988: 190).

The change in visual perception has nothing to do with "assumptions", but instead with changing where you fixate. And this is just an overt, bodily shift in spatial attention; it implies nothing about cognitive penetration.

This rebuttal—which I and others have called the *attention-shift interpretation*—says that the perceptual-cognitive situation has this rough causal structure:

(a) Cognitive state → Attention-shift → Perceptual experience

Pylyshyn (1999) and, more recently, Firestone and Scholl (2016) echo this analysis. The conclusion is that phenomena that fit schema (a) are relatively trivial to the architecture and epistemology of perception. They don't force important revisions to a modular theory, nor do they threaten the epistemic status of perception.

I grant to the modularist that phenomena that fit causal schema (a) will not count as cognitive penetration of perception. I then argue that there is an importantly distinctive causal schema (b), and that there are phenomena that fit this schema.

(b) Cognitive state → Non-agential selective attention → Perceptual experience

I then argue, both by appeal to consequences and by appeal to direct cognitive influence on perception, that those phenomena are best explained as cognitive penetration. Wayne Wu suggests that we should not so quickly accept that all instances of (a) are *not* cognitive penetration. I agree. But I do nonetheless make the concession that Wu charges me with. Here is why. In some ways, I grant the modularist's claims about the attention-shift interpretation and its applicability for sake of argument. And I do think the interpretation disarms appeals to bistable figures and the like, made by Churchland in those early debates. More substantively, I make the concession and focus on plausible cases of (b) for the following reasons.

First, interpretations that employ schema (a) make or rely upon questionable assumptions. They assume that attention is a spatial spotlight. They assume that attention is an action. They assume that attention is a gatekeeper between cognition and perception. By contrast, emphasis on schema (b) makes no such assumptions and indeed provides reason for doubting them. It thus underlines the varieties of attention (to which Wu is no stranger, see Wu 2011, 2013, 2014, 2017) and how selective attention is best understood as integrated with or part of vision.

Second, some invoke schema (a) and the attention-shift interpretation to dismiss the relevance of cases of perceptual expertise, claiming that experts just "know where to look" (Pylyshyn 1999: 360). Therefore, such cases are not cases of cognitive penetration or otherwise important cognitive influence on perception. But this explains only some cases of perceptual expertise. Emphasis on schema (b) better explains many cases of perceptual expertise. Feature-based and object-based selective attention is tuned to task and diagnostic needs, "weighted" according to the expert's background and cognitive training, in turn affecting perceptual salience and organization.

So while Wu may be right to resist the concession to the modularist that all cases of (a) are *not* cognitive penetration, I make that concession since I think it is emphasis on schema (b) that reveals these important lessons about attention, perception, and expertise.

Wu's second point is an important methodological one. We should not so quickly abandon emphasis on cognitive (im)penetrability, since such doctrines provide a useful bridge between behavioral and neural data in thinking about cognition-perception relations. And this is because there are competing, alternative explanations and the data underdetermines the choice between those explanations. Here again, I (mostly) agree. I'm explicit about the challenges from underdetermination in the book, and thus the argumentative structure is dominantly abductive. But Wu is on point: I rely almost entirely upon behavioral data and neural-physiological data, and I attempt to build the case that the convergence of this kind of data is best explained in terms of important, consequential cognitive influence on perception.

And here is a difficult balance I attempt to strike in the book: I continue to participate in the cognitive penetration debate in the first half and attempt to largely move away from it in the second half. Part of the reservation is inductive: it's not clear that this computational doctrine has gotten us very far (but Wu may have just provided a new counterexample). In any case, I see the importance of a computational approach, not least because an approach that employs all of these distinct levels of explanation may be optimal. Indeed, Wu's work on these scores has been groundbreaking, or should we say "floor shaking" (see Wu 2013, 2017).

And in his commentary, Wu offers a proposed case that makes good on his prescription. Drawing on his 2017 piece, Wu identifies a plausible instance where task-specific, instructed intention or expectation biases visual selective attention and thereby biases visual processing. Importantly, the computational model extends a plausible account about the neural correlates of selective attention (in terms of competition). And the behavioral evidence suggests that these cognitive effects on vision are *consequential* for the subject's behavior. The subjects are thus "effective visual agents".

This sounds like progress in the debate and it engages both old and new hands. And note further that if Wu is correct about his analysis of this case, we are talking about a pervasive phenomenon, and one that involves perceptual improvement. These are two themes that I stress in the book in the attempt to move away from idiosyncratic cases and apparently epistemically problematic cases of cognitive influence on perception.

To motivate the claim that imagination can be skillful, Amy Kind deflects two sceptical positions: the *no practice objection* and the *no feedback objection*. Against these objections, Kind co-opts research on acquisition of perceptual expertise to reveal the kind of practice and feedback that plausibly enables acquisition of imaginative skill. She highlights features of the first—repetition, self-initiated practice, self-feedback—and suggests that we will find analogs for these in the second. Agreed. And the case against the sceptic can be made stronger in at least two ways.

First—and this departs from the focus on *Thinking and Perceiving*—we should look to fiction and its appreciation. Engagements with fictions provide ample opportunity for both practice and feedback. Consider activities of appreciation and criti-

cism. The modes of engagement and disagreement here are partly explained by some individuals being more or less practiced—as well as more or less knowledgeable—in “playing the game of make believe” (Walton, 1990). And playing this game better, as it were, motivates more earnest practice which, in turn, can make one more imaginatively skillful. And it’s this kind of social context—arguing about a rich plot line, discussing the interpretation of a complex film—that provides feedback. So, fictions and their consumption provide feedback that improves imaginative skill. In one related empirical study on pretense in children, researchers tested children in closely matched pretend scenarios and counterfactual scenarios, and found that their performance success on those two tasks was significantly correlated (Buchsbaum et al., 2012).

Another (non-exclusive) way to resist the sceptic takes a different approach. Both the no practice and no feedback objections and replies focus on the preparation for (or acquisition of) the alleged imaginative skill. But another angle focuses on the outcomes of alleged imaginative skill. This is a performance-based rather than acquisition-based defence against the sceptic. Here we ask what evidence we have, in the form of individual differences, for performance of imaginative skill. Anecdote suggests that such differences exist, and so do imagery studies. There are plenty of established methods (and too many studies to cite here) for testing vividness of imagery, accuracy relative to some target or task, mental rotation and other manipulations in imagery. Individual differences in performance on these kinds of tasks might be explicable only by appeal to differences in imaginative skill (and therefore evidence for imaginative skill). Finally, here too Kind’s analogy with research on perceptual expertise will bear fruit.

Some of the research on perceptual expertise focuses on already established experts, where expert-level classification is determined by a standardized performance threshold. This is true for radiologists, bird and car experts, fingerprint examiners, elite athletes, and many other domains. Additionally, various tests aim to interfere with that level of performance, to identify its scope, to identify certain visual strategies used. And so Kind’s co-opting strategy will plausibly work here too, where some of the same methods of study might help to identify individuals with exceptional imaginative skill. One place to look for such analogues is in imagery-perception interference studies.

My second reply is a bit more critical but still very much sympathetic. Kind also co-opts some of the analysis of perceptual expertise to ground a claim about the epistemic value of imagination and imaginative skill. The thought goes like this. In the book, I suggest that at various places, operative in arguments for modularity, is an assumption that cognitive effects on perception will generally be epistemically problematic. Cases of expertise give us good reason to doubt this *pernicious cognitive effects assumption* (PCE) and therefore arguments that depend upon it. Kind suggests that this lesson applies analogously to a defense of the epistemic benefit of imaginative skill. An epistemic risk is supposed to attach to the fact that what we imagine is partly determined by cognition: our beliefs, desires, intentions, task selection, and so on. But if we reject PCE, this cognitive influence on imagination should, by analogy to cases of perceptual expertise, be epistemically innocuous if not beneficial.

Here is a worry about this analysis. Scepticism about the epistemic value of imagination might be grounded in the observation that imagination is voluntary *because* it's cognitively influenced; or it might be grounded in the observation that imagination is voluntary *because* it is non truth-bound. It may be the second that presents the real challenge. Suspicion about the epistemic benefit of imagination is grounded less in a concern about its being influenced by us (thus, our cognition), and more in a concern about it's not being sufficiently influenced or bound by, truth or reality. Indeed, it is plausibly voluntary *because* it is not bound to accurately representing the actual world. With perception, we start with the assumption that its function is world-responsive. Importantly, the claim that cognition can influence perception is not that the expert sees things that aren't there. She sees what is there differently (and for the better). But imagination, the sceptic might insist, is not so constrained and is therefore voluntary in an epistemically worrisome way.

So the better strategy to deflect the relevant sceptic may then be to lean on how the task or domain can provide constraints analogous to the worldly constraints on perception. Consider, once more, acquisition. If Kind is right, there is good reason to think that imagination can become skillful through practice and feedback. It can be sensitive to constraints on how practice proceeds (both in engagement with fictions and in counterfactual reasoning), and to feedback that underlines relevant constraints and when they have been violated. In a domain of perceptual expertise, be it forensics or football, there are norms and standards for performance success. Why should domains of imaginative skill be different? The social context, the selected task, the epistemic community, all can provide constraints on imaginative success and, accordingly, epistemic benefit.

Finally, Kind suggests that to further support my claim that some cases of perceptual expertise are genuinely perceptual, we should add consideration of imaginative skill. I'm receptive to the suggestion, and I want to identify one more bit of evidence that supports it. The basic thought here is that perceptual experts may be better imaginers/image-ers. Think of exceptional performance success, say where an athlete makes a novel or unexpected move. Such success is perhaps sometimes explained by invoking skilled imagery on the part of the athlete.

There is relevant empirical evidence here. Researchers have long identified perception-imagery *interference* (hindering or facilitating) (Farah, 1989; Craver-Lemley & Aterberry 2001; Pearson et al., 2008). And recent evidence suggests both behavioural and neural overlap between visual working memory (VSTM) and visual imagery (as well as visual perception) (Albers et al., 2013; Keogh & Pearson, 2011; Tong, 2013). Without space to detail this empirical work, the thought is this. If some experts are perceiving better and some of those freed neural resources serve as a workspace for both imagery and visual working memory, then Kind's claims about leveraging the case for imaginative skill by appeal to research on perceptual expertise would be plausible, and plausible partly because perceptual experts may also be good candidates for individuals with high imaginative skill within their domains of specialization.

Casey O'Callaghan has incisively forced me to think about some of the implications that come late in the book, concerning perceptual content and objectivity. A bit

of background on how I came to some of the relevant ideas. An epistemic argument for modularity goes as follows.

- (1) Perceptual representation can be largely accurate only if perceptual systems are informationally encapsulated.
- (2) Perceptual representation is largely accurate.
- (C) Therefore, perceptual systems are informationally encapsulated (ergo, modular).

The operative premise is the conditional in (1), and the modularist attempts to motivate it in various ways. They offer arguments that link reliability to encapsulation; in the book I argue that those arguments are too weak to secure (1). They sometimes appeal to the pernicious cognitive effects assumption, but that lacks supporting argument and looks dubious in the face of perceptual expertise studies. What's left? Another assumption, one that links reliable accuracy to robust objectivity. The accuracy of perceptual content is determined by, or just is, purely mind-independent Objectivity\*. Standards for accuracy – what makes a perceptual experience accurate – are given entirely by external reality, by the stimulus, absent any subjective influence or determinant. I argue that Objectivity\* is the wrong notion for perceptual accuracy. O'Callaghan has challenged me to clarify what I'm committed to in making that claim.

O'Callaghan offers three distinctions. He writes “Universality and objectivity are orthogonal to each other. Traditional modularists overestimated universality in perception. But universality just makes it easier practically and epistemically to resolve disputes. Objectivity does not require it” (P. 9). First, distinguish pure mind-independent *Objectivity\** from *Inter-subjective objectivity*. Commitment to the second, at least in some radical form, is supposed by the modularist to generate a worrisome subjectivism, one that results in scepticism. Second, distinguish *Content universalism* from *Content pluralism*. Commitment to the second is supposed to generate a radical relativism, one that doesn't allow for dispute resolution. The importance of these distinctions comes into focus with a final distinction. O'Callaghan highlights “the importance of distinguishing an account of how accuracy conditions and contents are determined from an account of the nature of accuracy conditions and contents themselves. The objectivity of each ought to be considered as a distinct question” (P. 9). Thus we distinguish the causal determinants of accuracy conditions and content from the constituents of accuracy conditions and content. These distinctions in place, O'Callaghan asks: How radical or revisionary is the view about perceptual content?

The short answer: Inter-subjective and Objective features of the world causally determine content *and* constitute the accuracy conditions for that content. This entails diversity of content, a content pluralism. But it is compatible with (but does not entail) that perceptual content is constituted by purely Objective\* features of the world. Now for the longer answer.

Perception is a process of *objectification* (Burge, 2010). When one has the visual experience as of something being pink or flat, this involves attribution of those features to some object outside of and independent of oneself. The instances of pinkness and or flatness are experienced as features not of one's subjective experience but of

objective reality: say, the glass of rosé or the marble countertop. But because perception is a capacity for sensory representation by *an individual*, its representations are context and perceiver-dependent. In this way, perception is an individual process that reflects egocentric features such as perspective but also reflects the roles that the individual subject, *qua* person, and sometimes her context, play in that process.

Recall the example (discussed by O’Callaghan) of two distinct experts watching penalty kick after penalty kick. The quality or success of perceptual representation here depends on what is behaviourally relevant in the context of perception. The two perceivers are members of, respectively, the community or domain of sports medicine and of playing football (soccer). Even when viewing the same events, by virtue of their distinctive domain-centred goals, two perceivers can perceive those events differently while both can be accurate and useful.

Does this yield an inter-subjective account of accuracy? The answer I tend to favor in the book is: yes. Whether a perceptual experience is accurate can be partly *determined* by the behaviourally relevant information in the context and, further, sometimes by the goals, needs, and standards specific to a domain of inquiry or community. This is a relativism about accuracy conditions, but checked by the epistemic community. I think it is what O’Callaghan calls “intersubjective epistemological objectivity”, and it admits theory-ladenness of perception. What about constituents of accuracy conditions? On the standard line, a specification of accuracy conditions is a specification of how the world would need to be for one’s experience to be accurate. The satisfiers of those conditions are, in the good cases, objects and features of the world. Those are the contents of one’s experience thus understood. So we can take the satisfiers to be objective mind-independent entities. But these kinds of conditions themselves, the ones that specify those satisfiers, are abstractions. They are norms or standards for value along some dimension. Accuracy is a success concept, but what counts as accurate—what makes an experience accurate or successful—can vary from individual to individual, and context to context.

Thus, accuracy conditions are sometimes partly *constituted* by behaviourally relevant information in the context and, further, sometimes by the goals, needs, and standards specific to a domain of inquiry or community. This is a subjective objectivism about accuracy conditions, but again checked by the epistemic community. It’s for these same reasons that the threat of scepticism here is disarmed. The conditions are not constituted by facts concerning a single individual, but instead derive from a specialized community, be it football or forensic examination.

One might insist that accuracy just cannot be relativized to perceiver or domain. Whether an experience is accurate is determined by whether the feature combinations selected by vision are in fact instanced in the visible environment; this remains a robustly objective matter. Perhaps, but this just reveals that accuracy does not exhaust the success of perception. And an emphasis on perceptual expertise suggests that perceptual success is a more interesting notion. Expert perceivers are successful along a variety of measures that align with aims for action, task performance, and worldly engagement. So even if perceptual sensitivity to patterns and organizational features, to task demands, and less susceptibility to distraction, do not determine perceptual *accuracy*, they do plausibly determine perceptual *success*. Perceptual success subsumes perceptual accuracy. And perceptual experts are often more successful not by

being more accurate but by being perceptually sensitive to features relevant to their task. This may imply some interesting revisions to how we think about content vis-à-vis accuracy, perhaps thinking about it instead in terms of certain kinds of success conditions. And, finally, this remains compatible with the safest answer regarding constituents of perceptual content: those constituents are just objective features of the world.<sup>1</sup>

As the book progresses, I favor malleability over modularity, success over accuracy, understanding over propositional knowledge, inter-subjectivity over objectivity. How much of us as individuals, then, is in what we perceive? A lot. The contents of experience are partly causally determined by who we are, what we know and have learned, and what we do well. And the conditions for perceptual success and accuracy—if not contents themselves—are partly constituted by, made what they are by, objective facts, but also by who we are, our tasks, and our sometimes very specialized domains of activity. That, anyway, is how I come to this conclusion: our contact with the world is very much our own.

**Acknowledgements** This symposium was presented at the APA Pacific Meeting in Vancouver, 2022; many thanks to all that were present at that event, and for their feedback. Thanks also to Colin Macleod for his editorship of this special issue.

## References

- Albers, A. M., Kok, P., Toni, I., Dijkerman, H. C., & de Lange, F. P. (2013). Shared Representations for Working Memory and Mental Imagery in early visual cortex. *Current Biology*, *23*, 1427–1431. <https://doi.org/10.1016/j.cub.2013.05.065>.
- Buchsbaum, D., Bridgers, S., Weisberg, D., & Gopnik, A. (2012). The power of possibility: Causal learning, counterfactual reasoning, and pretend play. *Phil Trans R Soc -Bio Sci*, *367*, 2202–2212.
- Burge, T. (2010). *Origins of Objectivity*. Oxford University Press.
- Churchland, P. M. (1988). Perceptual plasticity and Theoretical Neutrality: A reply to Jerry Fodor. *Philosophy of Science*, *55*, 167–187. <https://doi.org/10.1086/289425>.
- Craver-Lemley, C., & Arterberry, M. E. (2001). Imagery-induced interference on a visual detection task. *Spatial Vision*, *14*, 101–119. <https://doi.org/10.1163/156856801300202887>.
- Farah, M. J. (1989). Mechanisms of imagery-perception interaction. *Journal Of Experimental Psychology: Human Perception And Performance*, *15*, 203–211. <https://doi.org/10.1037//0096-1523.15.2.203>.
- Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for ‘Top-Down’ Effects. *Behavioral and Brain Sciences*, *39*, 1–72. <https://doi.org/10.1017/S0140525X15000965>.
- Fodor, J. A. (1988). A reply to Churchland’s ‘Perceptual Plasticity and Theoretical Neutrality’. *Philosophy of Science*, *55*, 188–198. <https://doi.org/10.1086/289426>.
- Keogh, R., & Pearson, J. (2011). Mental imagery and visual working memory. *PLoS One*, *6*(12), e29221. <https://doi.org/10.1371/journal.pone.0029221>.
- Pearson, Pearson, J., Clifford, C. W., & Tong, F. (2008). The functional impact of mental imagery on conscious perception. *Curr. Biol.*, *18* (2008), pp. 982–986.

<sup>1</sup> This safe answer is compatible with the account I’ve given but, again, not entailed or required by it. And it’s worth noting that what one says here depends partly on one’s independent theory of the structure of perceptual content. And there are defensible views, for example indexical content views, that do include some index to the perceiver as constituents of content.



- Pylyshyn, Z. (1999). Is vision continuous with cognition?: The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, 22(3), 341–365. <https://doi.org/10.1017/s0140525x99002022>.
- Tong, F. (2013). Imagery and visual working memory: One and the same? *Trends in Cognitive Sciences*, 17(10), 489–490.
- Walton, K. (1990). *Mimesis as make-believe*. Cambridge, MA: Harvard University Press.
- Wu, W. (2011). What is conscious. Attention? *Philosophy and Phenomenological Research* 82, 93–120. <https://doi.org/10.1111/j.1933-1592.2010.00457.x>.
- Wu, W. (2013). Visual spatial constancy and modularity: Does Intention Penetrate Vision? *Philosophical Studies*, 165, 647–669. <https://doi.org/10.1007/s11098-012-9971-y>.
- Wu, W. (2014). *Attention*. London: Routledge.
- Wu, W. (2017). *Shaking up the mind's ground floor*. The Cognitive Penetration of Visual Attention. *Journal of Philosophy* 114, 5–32. <https://doi.org/10.5840/jphil201711411>.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.