

UC San Diego

UC San Diego Previously Published Works

Title

Perceptual representation, veridicality, and the interface theory of perception

Permalink

<https://escholarship.org/uc/item/838271sv>

Author

Cohen, Jonathan

Publication Date

2014-10-01

DOI

10.3758/s13423-014-0782-3

Peer reviewed

Perceptual representation, veridicality, and the interface theory of perception

Jonathan Cohen*

He who would distinguish the true from the false must have an adequate idea of what is true and false.

— Spinoza, *Ethics* II, proposition 42: proof.

Hoffman, Singh, and Prakash (henceforth, HSP) are not fans of veridical perceptual representation. They argue that perception was not selected for veridical representation, hence that, contrary to a very widespread consensus, there's much less of the latter than you might expect in perception. Moreover, they use this conclusion to argue not only that perception is ordinarily not a veridical source of information about the way the world is, but that standard Bayesian approaches to computational vision are seriously misformulated, that the objects of perception are not public mind-independent objects, and for various other self-consciously iconoclastic conclusions. And they put forward an alternative "interface" theory, on which perception is a kind of adaptively useful but truth-obscuring veil mediating between perceiver and perceived. Pretty heady — and radical — stuff, to be sure, as HSP are well aware.¹

Unfortunately, the arguments HSP offer us against the veridicality of perception depend crucially on a serious misapprehension in the early going about what veridicality amounts to. This mistake both undercuts their arguments against perceptual veridicality and prevents them from seeing that their own preferred conception of perception is itself committed to veridical representation, rather than an alternative to it. Moreover, since HSP's further radical conclusions depend on their conclusion that perception fails to represent veridically, the failure of their arguments for the latter also undermines the former.

In what follows I'll begin by explaining how HSP's misunderstanding of veridicality prevents them from asking the questions they want to ask (§1), and then going on to note that we can reformulate their questions more fruitfully by appeal to the notion of content (§2). Alas, I'll next point out, given this reformulation, HSP's case against perceptual veridicality — both their motivating examples and their evolutionary games — depends on unsupported assumptions about the content of perceptual states that

*Department of Philosophy, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0119, joncohenundo@aardvark.ucsd.edu

¹HSP embrace their iconoclasm, proudly locating themselves in a tradition of brave rebels stretching back to Kepler, Copernicus, Galileo, and Bruno (ms31).

we have reason to reject (§§3–4). Finally, I’ll show that, while HSP’s simulations can be used to raise a challenge (though not the one they intend) to widespread appeals to perceptual veridicality, this challenge, too, is unconvincing (§5). In the end, I’ll conclude, HSP give us no reasons to abandon the standard view that perception veridically represents the world.

1 Veridicality mischaracterized

As HSP point out, it is a standard part of contemporary thinking about the mind that, by and large and with many individual exceptions, perception functions to represent veridically aspects of the distal world. They want to convince us, however, that this standard view is false.²

HSP frame the question of veridicality in terms of their notion of a “perceptual strategy” — an “onto measurable function $P : W \rightarrow X$, where W denotes a measurable space of states of the world and X denotes a measurable space of perceptual experiences” (ms9).³ They formulate several, increasingly weak, types of perceptual strategies, ranging from “naive realism” ($X = W$; P is an isomorphism) to “strict interface” (“no perceptions are veridical ($X \not\subseteq W$) and no structure of W is preserved other than measurable structure (P is not a homomorphism)”) (ms14). And then they proceed to argue, on evolutionary game theoretic and other grounds, that perception is best characterized in terms of the weakest sort of perceptual strategies, on which perception does not veridically represent the world.

But we should notice immediately that this formulation of the options commits the proponent of veridical representation to something deeply implausible, and so stacks the deck against such views before we even get to the arguments. The condition that HSP here take as defining veridical perception — that $X \subseteq W$ — amounts to an absurd identity theory, on which x veridically represents y iff $x = y$. But, exactly because such an identity theory is absurd, it is an unreasonable assumption for HSP to foist on their opponents.⁴

Of course we should deny that $X \subseteq W$. For one thing, setting aside those (presumably statistically vanishing) cases where we perceive our own mental states,

²Needless to say, proponents of veridical representation in perception allow that many — perhaps a statistical majority of, or even an overwhelming statistical majority of — individual perceptual representations are non-veridical. The claim that proponents of veridical perceptual representation hold, and that HSP want to argue against, is that perception crucially involves a capacity (normatively) for representing the way the world is in some respects, and that exercises of this capacity count as veridical when the world is that way. Given that this is the claim at issue, it won’t suffice to argue against it merely to point out that there are (even quite a lot of) individual perceptual states that represent falsidically. (This is a clarification, not a criticism: HSP themselves appear to be aware of this point, since they argue against the claim on different grounds.)

³This is their definition 1, which ignores dispersion, rather than their definition 2, which does not; nothing that follows hangs on this simplification.

⁴Indeed, I’m not sure that even HSP themselves believe that veridical representation entails the identity theory. At various points they appear to shrink back from that condition, and instead talk about successful veridical representation relations as instances of “isomorphism” (ms7), “homomorphism” (ms21, ms22), or cases in which the states “resemble” (ms16, ms72) or “depict” (ms3, ms16) the world. All of these are *prima facie* more plausible as accounts of veridical representation (though certainly not without difficulties of their own), but all are weaker than (a fortiori, different from) the identity relation. Unfortunately, they don’t say more about how they are understanding these weaker relations, or whether by invoking them they mean to be softening their pessimism about veridical representation.

elements of X and elements of W fail to share, *inter alia*, locations (in the head, in the world, respectively), mind-independence (elements of X are mind-dependent, elements of W are mind-independent), survival conditions (elements of X pass out of existence when I sleep/die, elements of W don't), and representational status (elements of X have it, elements of W don't). It is no less uncontroversial (barring some kind of anachronistic Aristotelian hylomorphism on which the perceiving mind takes on the form of perceived items) that elements of X and corresponding elements of W differ in more or less all of their other properties: when a perceptual experience $x \in X$ represents the world state $w \in W$ in which a red tomato sits on a rectangular table, x does not itself exemplify any of the properties (e.g., *redness*, *being a tomato*, *sitting upon*, *green*, *rectangularity*, *being a table*) that constitute w . Again, given those differences between their elements, *of course* it can't be that $X \subseteq W$.⁵ Moreover, we should note that the identity theory is (obviously) untrue about garden variety non-perceptual cases of veridical representation. For example, the English words 'snow is white' (/the French words 'la neige est blanche'/etc.) veridically represent that snow is white, but surely are not identical with (and have many different properties from) snow, whiteness, or the proposition that snow is white. The photograph may represent veridically that Grandma's sweater is blue, but does so without itself being identical to (or sharing properties with) Grandma, the sweater, blueness, or the proposition/state of Grandma's sweater's being blue. And the map may represent veridically that the river is east of the mountains, but does so without being identical to (or sharing properties with) the river, the mountains, the relation of being east, or the proposition/state of the river's being east of the mountains.

In short, the understanding of veridical representation HSP assume seems deeply implausible (and, for what this is worth, not accepted by any serious theorist I am aware of), so shouldn't be built into the meaning of the standard claim that perception veridically represents the world.⁶

2 Veridicality revisited: Content

While HSP's official understanding of veridicality is flawed, there is an improvement in the neighborhood that allows us to frame their question more fruitfully. What is needed is the (standard) idea that perceptual states have *content* — intuitively, what they carry information about, tell us about, or say about, the world, and that can be evaluated for truth or falsity (the world either is as they say, in which case the states are veridical, or is not as they say, in which case the states are not veridical).⁷

⁵Additionally, if somehow, members of W were identical to members of X , that would make perception awfully mysterious: it's hard to imagine what causal mechanism could make it the case that distal world states (outside of heads) literally are states of cognitive systems (inside heads).

⁶Obviously, I don't mean to suggest that standard views can't be rejected, but only that one would like to see motivation for rejecting them — especially if rejecting them has serious costs for those urging the rejection, as we've seen is the case here for HSP.

⁷It's worth mentioning that there are indeed dissenters from the standard appeal to content, including Brewer (2006); Campbell (2002); Martin (2006); Travis (2004), as well as dissenters from the dissent, such as Burge (2005); Byrne (2002, 2009); Dretske (1995); Harman (1990); Pautz (2010); Siegel (2010); Tye (1995). But HSP wouldn't want to accept the no-content view of Travis et. al. as a way out of the problems raised in §1, since (precisely because it forswears the notion of content) this view entails, pace HSP's desired conclusion, that perceptual states are just as incapable of misrepresenting the world as they are of representing it veridically. (Terminological caution: The

Given this modest (and standard) additional apparatus, we can now say that a perceptual state x veridically represents a world state y iff (i) x has the content y , and (ii) y obtains.⁸ This puts us in a position to redefine HSP’s perceptual strategies as functions from the space of world states W to (not the space of perceptual experiences, X , but) the space of *contents* of perceptual experiences C_X . And we can reformulate their question about veridicality this way: should we think about perception in terms of perceptual strategies that represent veridically (i.e., perceptual experiences have contents, or say things about the world, and the world really is as those experiences say it is) or perceptual strategies that represent the world falsidically (i.e., perceptual experiences have contents, or say things about the world, but the world is not as those experiences say it is)?

3 Content and Falsity

It would seem, then, that there are many theoretical advantages, both in general and for HSP’s argumentative purposes in particular, for appealing to the notion of perceptual content in these discussions. That’s the standard view anyway, and as we’ve been seeing, it is well-motivated, both by the need to avoid the problems considered in §1 and by the apparent truism that you can’t say whether something is veridical or not without first knowing what it is saying.

Alas, this truism turns out to be somewhat embarrassing for HSP, in so far as the examples they use to motivate their denial of perceptual veridicality depend crucially on controversial (and unsupported) assumptions about just what contents states have. Thus, they offer (ms6) as an instance of non-veridical perception the case of dragonflies that locate suitably watery positions for laying eggs by detecting horizontally polarized light reflections, so sometimes disastrously end up laying eggs in locations that reflect horizontally polarized light but won’t support dragonfly eggs (e.g., oil slicks and shiny tombstones). In the same vein, they point (ms6) to the example of male jewel beetles that locate sexual partners by detecting glossy, dimpled, and brown surfaces normally found on the wing-casings of female jewel beetles, but sometimes end up trying unsuccessfully to mate with objects that possess the right sort of glossy, dimpled, and brown surfaces but are regrettably unhelpful with the reproductive task at hand (e.g., discarded empty beer bottles).

Unfortunately, though HSP intend these as clear examples of non-veridical perception (which they hope will soften us up towards a more widespread skepticism about perceptual veridicality), it’s highly unobvious that that’s the right description. Much depends on how we understand the contents of the relevant states. Admittedly, if we suppose that the relevant dragonfly perceptual states represent the content *water at location l* or *location l suitable for laying eggs* (and that dragonfly visual systems use the optical cue of horizontally polarized light reflection merely as a trigger for tokening that content), then, so long as l is the location of an egg-inhospitable oil slick, the dragonflies’ perceptual states are falsidical: they represent the world as being a way that it is not. But if, on the other hand, dragonfly perceptual states represent the content *horizontally polarized light at location l* (and dragonflies have evolved the response of laying eggs

no-content view is often discussed under the label “naïve realism,” though discussants clearly mean something different by it from what HSP mean by the term, as they note on ms10).

⁸Of course, all the key notions in this sort of account are themselves up for grabs. In particular, and among other things, this kind of generic formulation leaves it open what contents are, what it takes for one thing to have another as content, and what it is for a content to be true of the world.

when they token that content), then these perceptual states say something *true* about the world, so are veridical. Mutatis mutandis: if jewel beetle perceptual states represent the content *b is an appropriate sexual partner*, they are falsidical; but if they represent the content *b displays a glossy, dimpled, brown surface*, they are veridical.⁹

The present point is not that we are forced to make content assignments on which the relevant states turn out to be veridical after all; it is that, pending some argumentative support, HSP aren't entitled to make content assignments on which they turn out to be non-veridical. Of course, even if these cases were counted as instances of non-veridical perception, that wouldn't succeed in motivating a more widespread skepticism about perceptual veridicality without additional reason to think that their lack of veridicality is typical and generalizable. HSP are aware of this, and present evolutionary arguments (about which more anon, see §4) in the hope of convincing us that veridicality is unlikely as a general matter, and so of filling in the generalizing step of their intended induction. But what we're seeing is that it's up in the air whether these supposedly parade examples really are instances of non-veridical perception, hence that HSP owe us an argument for the content assignments they take for granted even to get their inductive case against perceptual veridicality started.

Much the same point applies to HSP's example of a computer interface, which they use to motivate/explicate their alternative ("interface") theory of perception. Specifically, it appears that their case for treating the computer interface as non-veridical rests on implausible (and unargued) assumptions about what it takes to have representational status, and about the content of the interface's representational states.

HSP appear to think it obvious that computer interfaces are not veridical representations:

Suppose there is a blue rectangular icon in the upper right corner of the desktop for a text file that you are editing. Does this mean that the text file itself is blue, rectangular or in the upper right corner of the laptop? Of course not. Anyone who thinks so misunderstands the purpose of the desktop interface. Its icons are not meant to resemble anything in the computer; they are not veridical representations. But they are intended to guide useful behaviors. If, for instance, you drag the blue icon to the trash you can delete the text file; if you drag it to the icon for an external drive, you can copy the file (ms15–16).

But the observations here do not show what HSP take them to show.

HSP are surely right that the text file lacks properties that the icon exemplifies. But as we saw in §1, no sane account of representation (viz., not one that could possibly apply to linguistic, depictive, or map-like cases, and not one that is even prima facie plausible about perception) demands identity or widespread property sharing between vehicle and content, so that's no reason to deny that there is representation going on here. If, as seems overwhelming natural, the icon *does* represent (as even HSP come perilously close to recognizing: "there is [an] icon ...for a text file..."), then is it

⁹You might think that there is a straightforward evolutionary preference for assigning contents like *location l suitable for laying eggs*/*b is an appropriate sexual partner* over contents like *horizontally polarized light at location l*/*b displays a glossy, dimpled, brown surface* because the former contents have obvious and direct adaptive significance, while the latter have at best instrumental adaptive significance. But that would be a mistake, even conceding (controversially) that adaptive significance fixes content. Evolution can't prefer either of these candidate contents over the other because, so long as the instruments are sufficiently well-correlated with the ends in the (now past) evolutionary environment in which selective pressures applied, selecting for the one will have been precisely as adaptive as selecting for the other (cf. Fodor 1990, p. 73).

veridical? As before, that depends on (i) its content and (ii) how the world is. If its content is *there is a blue rectangular text file in the upper right corner of the laptop*, then it represents falsidically: for there is no blue rectangular text file in the upper right corner of the laptop. But, exactly as HSP correctly note, that would be a crazy understanding of the icon's content. Icons, like other representational artifacts, plausibly have as their contents what their users/makers intend for them to represent; and what users/makers of this icon intend for it to represent is presumably something like *there is a text file stored on a portion of the machine's hard drive that has been assigned (by its users/makers) the label 'Desktop'*. And there is a text file on a portion of the machine's hard drive that has been assigned (by its users/makers) the label 'Desktop'; hence, on the latter content assignment, the icon represents veridically.¹⁰

If we follow HSP, then, in thinking that computer interfaces are useful models for understanding perception, then what they have said so far gives us no reason for denying that perception veridically represents the world, and every reason for affirming that it does.

4 Evolutionary games

Though HSP offer the examples discussed above in the hopes of smoothing the way for skepticism about perceptual veridicality, their evolutionary game theoretic argument forms the heart of their case. They want to argue that, as a general matter, perceptual strategies that don't represent veridically ("interface strategies") outcompete those that do under selective pressure, hence that our sort of perceptual systems, shaped as they are by natural selection, must employ strategies of the former type.

Thus, they compare (ms20–24) a strategy that represents the quantity of resources in a territory against a strategy that is "tuned to payoffs" (i.e., conveys information about that very quantity rewarded by the selection function) in the same territory, and find that, in Monte Carlo simulations over many iterated versions of the game, selection strongly favors the latter over the former. Or, in their words, "The key insight from these evolutionary games is this: Natural selection tunes perception to payoffs, not to truth" (ms23), except in the measure-zero subset of cases where payoffs and truth vary monotonically.

Additionally, HSP appeal to evolutionary games using genetic algorithms to argue that perceptual strategies involving veridical representation shouldn't even be expected

¹⁰Surprisingly, HSP appear to accept this point:

One might for instance argue that — contrary to the "simplify and hide" hallmark of interfaces that we have touted — folder icons on a desktop do faithfully depict the details of real folder hierarchies in the computer. This critique is well taken. However, it is the interface strategy as precisely defined, not as metaphor, which lives or dies by the sword in our evolutionary games (ms16).

I'm afraid I don't see how HSP can so blithely concede the point, treating the case at issue as a mere imperfect metaphor, given that the very same objection arises in precisely the same form in other cases which (if the talk about living and dying by swords is to be believed) they are not prepared to give up. Thus, as we have already seen in considering the dragonfly and jewel beetle examples, exactly the same worry applies in perceptual cases, where HSP's claim that states represents falsidically fails without their controversial and unsupported assignments of content to those states. Moreover, as we'll see in §4, precisely the same issue arises for the states in HSP's evolutionary games. If the objection is damning with respect to the state of the computer, as they are apparently prepared to concede, then the identical objection is no less damning for their central claims about perceptual representation.

to emerge at all under natural selection. In the game they describe (ms24–28), a robot forages for soda cans on a 10x10 grid where each square has 0–10 cans. Robots token one of two mutually exclusive perceptual states, red or green, whenever there are n cans in the currently perceived square, for every n between 0 and 10. They have “genes” specifying both a perceptual strategy (a mapping from numbers of perceived cans to perceptual states), and an action plan (a mapping from possible perceptual states to possible actions). The fitness function is specified by the sequence $\langle 0, 1, 3, 6, 9, 10, 9, 6, 3, 1, 0 \rangle$, where the value in the i th position is the payoff for grabbing in a square containing i cans. Fitness is computed over many runs of a strategy, and the fittest strategies are allowed to pass on their genes in a way that mimics sexual reproduction with mutation. After 500 generations, the strategies that emerge are those in which robots assign one perceptual state to low payoff squares (those having 0, 1, 9, or 10 cans) and the other to high payoff squares (those having any other number of cans). What does not emerge, according to HSP, is a realist strategy tuned to truth:

A strategy tuned to truth would see squares having between 0 and 5 cans as, say, red and squares having between 6 and 10 cans as green, so that the perceived color would be as informative as possible about the true number of cans. But such a realist strategy would provide no information about payoffs (since red squares would have the same expected payoff as green squares) and would thus fail to guide effective foraging (ms27).

Hence, they conclude, we shouldn’t expect that veridicality-involving perceptual strategies would ever emerge as phenotypes: “it’s likely that veridical strategies never enter the playing field. They’re so unfit that they’re not worth trying” (ms28).

But the results HSP describe do not show anything of the sort. For, advertising to the contrary notwithstanding, it’s not true that the evolutionary simulations under discussion compare realist, veridicality-requiring perceptual strategies against interface strategies that do not require veridicality. Rather, these games compare one veridicality-requiring strategy against a different veridicality-requiring strategy.

HSP are right that one of the perceptual strategies they compare in their first game (their “realist” strategy) veridically represents the distribution of resources: the perceptual states here carry information about the quantity of resources in each territorial region, and the regions have exactly the quantities represented by those perceptual states, so the states are veridical. But the alternative (“interface”) strategy they pit against their realist strategy is no less committed to veridical representation. Of course, the states in the alternative strategy don’t represent/carry information about/tell us about quantity of resources in each region, so they can’t have resource quantities as their content. But these states do (by construction) represent/carry information about/tell us about the payoffs available in each region — which is just to say that those perceptual states have territorial payoffs as their contents. And those states are (again by HSP’s own stipulation) *veridical* representers of the world: the payoff in a region just is exactly what the perceptual state represents it as being.¹¹ It’s true that the HSP-preferred strategy veridically represents something different about the world from what the HSP-dispreferred strategy veridically represents. But both are, in exactly the same ways and to exactly the same extent, strategies in which perception veridically represents the world.

¹¹In fact, HSP never consider strategies that represent but do so non-veridically. However it is obvious that, *ceteris paribus*, a strategy representing payoffs won’t incur an adaptive advantage if it represents *non*-veridically, so it’s not open to HSP, who want to infer from adaptive advantage to the structure of perceptual strategies, to hold that states represent but do so falsidically.

The same remark applies to HSP's other game. Again, they are right that the losing perceptual strategy in their comparison is committed to veridical representation: it veridically represents (with low precision) the number of cans in a cell. But, also again, they are wrong to think that the winning strategy in their comparison avoids veridical representation. On its face, the winning perceptual strategy is one that represents/carries information about/tells us about whether the current cell is one in which grabbing has a high payoff. And it does so veridically (lest it fail to confer a selective advantage after all).

5 Veridically representing payoffs and veridically representing other things

Because, as we have seen, HSP's simulations only ever involve competitions between perceptual strategies that represent veridically, their simulations are not in a position to answer either of the following questions:

- (1) Does a perceptual strategy that represents the world confer a selective advantage relative to a perceptual strategy that does not represent the world?
- (2) Does a perceptual strategy that represents veridically confer a selective advantage relative to a perceptual strategy that represents non-veridically?

The question the simulations are potentially in a position to answer is, rather,

- (3) Does a perceptual strategy that represents veridically one sort of content about the world — viz., content about the distribution of payoffs — confer a selective advantage over a perceptual strategy that represents veridically a different sort of content about the world — viz., content about the distribution of factors on which payoffs depend?

Question (3) is, of course, not the one HSP intended to ask. And the answer to it provided by their simulations (selection appears to prefer veridical representation of the distribution of payoffs to the veridical representation of the distribution of other quantities) is deeply unsurprising. Payoffs are, after all, just what selection rewards; and it is to be expected that (*ceteris paribus*) representing the distribution of payoffs gets creatures closer to securing those payoffs (hence, to enjoying a selective advantage) than does representing the distribution of anything that comes apart from payoffs.

Still, this unsurprising answer to a question HSP didn't ask might suggest an HSP-like way of arguing for a conclusion in the vicinity of their intended view. Namely, one might argue that, since veridical representation of the distribution of payoffs outcompetes veridical representation of anything else under simulation, evolution must have shaped perception to represent veridically only the distribution of payoffs. This conclusion would put significant limits on perceptual veridicality (even if it did not license HSP's completely general skepticism about the notion), and would indeed threaten the standard appeals to veridical perceptual representation (say, of color, form, motion, etc.) in cognitive science that they are so keen to reject.

This argument might be persuasive if it were true that the traits it pits against one another (the capacity to represent veridically the distribution of payoffs, the capacity to represent veridically the distribution of other features) were independent, hence subject to independent selective pressures. If organisms really could have only one of these traits, one would indeed expect that natural selection would favor those with the first. The simulations HSP discuss treat the traits as independent in just this way, so

encourage the view that veridical representation of payoff distributions is completely unlinked to the veridical representation of other features. But this assumption is utterly unrealistic for the kinds of perceivers we know about. In such perceivers, the capacity for perceptual representation of payoffs is strongly positively correlated with the capacity for perceptual representation of other features: indeed, to the extent that we perceptually represent the distribution of payoffs, we do so (partly) *because* we represent the distribution of color, form, and all the rest. (Recall that, for example, HSP's dragonflies track reproductive payoffs by tracking horizontally polarized light reflections, that their male jewel beetles track payoffs by tracking certain kinds of glossy, dimpled, and brown surfaces, etc. This situation is, it would seem, entirely typical.) But because the two traits at issue are highly correlated in the evolutionary relevant environment, selection for one of them just is selection for the other. As such, even if we take the simulations to show that evolution must have shaped perception to represent veridically the distribution of payoffs, they cannot show that it shaped perception to represent veridically *only* the distribution of payoffs.¹²

6 Conclusion

I suppose there's a way of seeing in the foregoing a vindication for HSP. For, if veridical perceptual representation means what they take it to mean, then they are right: there isn't much of it in perception.

But this victory turns out to be rather hollow. For veridical perceptual representation can't be reasonably understood in HSP's way, and for reasons you don't need any of the arguments they offer us to see. Nor is there any reason to believe that the many proponents of the notion are working with anything like HSP's understanding, so are in need of correction on this point. Moreover, a more reasonable reading of the claim that perception represents the world is left untouched by HSP's considerations (both their supposedly motivating examples and their evolutionary game theoretic arguments). Finally, it would seem that HSP's supposedly alternative, "interface" picture of perception is in fact no less committed to veridical perceptual representation than the views they aim for it to replace.

As far as I can tell, then, HSP have given us no reason for doubting that perception veridically represents the world.¹³

¹²There are other, more generic, worries about the external validity of HSP's evolutionary games, which leave unmodeled a variety of real-world constraints that can impact the course of actual phenotypic expression dramatically. For example, they say nothing about how physical constraints shape the space of options, and implicitly assume that the emergence of the representational capacities at issue is independent of (imposes no costs on) any other possible or actual selectively advantageous traits of the organism. (Compare: Perhaps human beings who can run at 100mph but are otherwise exactly like us would outcompete under simulation human beings who cannot run at 100mph and are otherwise exactly like us. It's probably not a good idea to infer from that fact to the conclusion that actual human beings must be able to run at 100mph.) It's not an objection against evolutionary simulations, per se, to note that they often fail to incorporate constraints of this kind; but it is a reason to be cautious about inferring from the dominance of a phenotype under simulation to the conclusion that that phenotype must be expressed in the actual world.

¹³I am grateful to Craig Callender, Matthew Fulkerson, Don MacLeod, and Brian Tracz for discussion of these issues.

References

- Brewer, Bill (2006), "Perception and Content," *European Journal of Philosophy*, 14, 2, pp. 165–181.
- Burge, Tyler (2005), "Disjunctivism and Perceptual Psychology," *Philosophical Topics*, 33, 1, pp. 1–78.
- Byrne, Alex (2002), "Intentionalism Defended," *The Philosophical Review*, 110, pp. 199–240.
- (2009), "Experience and Content," *Philosophical Quarterly*, 59, 236, pp. 429–451.
- Campbell, John (2002), *Reference and Consciousness*, Oxford University Press.
- Dretske, Fred (1995), *Naturalizing the Mind*, Originally delivered as the 1994 Jean Nicod Lectures, MIT Press, Cambridge, Massachusetts.
- Fodor, Jerry A. (1990), "A Theory of Content, II: The Theory," in *A Theory of Content and Other Essays*, ed. by Jerry Fodor, MIT Press, Cambridge, Massachusetts, pp. 89–136.
- Harman, Gilbert (1990), "The Intrinsic Quality of Experience," in *Philosophical Perspectives: Action Theory and Philosophy of Mind*, ed. by J. Tomberlin, 4, Ridgeview Publishing Company, Atascadero, California, pp. 31–52.
- Martin, Michael G. F. (2006), "On Being Alienated," in *Perceptual Experience*, ed. by Tamar S. Gendler and John Hawthorne, Oxford University Press.
- Pautz, Adam (2010), "An Argument for the Intentional View of Visual Experience," in *Perceiving the World*, ed. by Bence Nanay, Oxford University Press, New York.
- Siegel, Susanna (2010), "Do Experiences Have Contents," in *Perceiving the World*, ed. by Bence Nanay, Oxford University Press, New York.
- Travis, Charles (2004), "The Silence of the Senses," *Mind*, 113, 449, pp. 57–94.
- Tye, Michael (1995), *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*, MIT Press, Cambridge, Massachusetts.