Analytic Philosophy

WILEY

ANALYTIC PHILOSOPHY

ORIGINAL ARTICLE

Constructivism for Consciousness

John Halpin

Department of Philosophy, Oakland University, Rochester, Michigan, USA

Correspondence: John Halpin (halpin@oakland.edu)

Received: 30 May 2025 | **Accepted:** 2 April 2026

Keywords: consciousness | physicalism | pragmatic best-system account

ABSTRACT

The goal of this essay is to promote a non-traditional physicalism for consciousness much as metaethical constructivisms are meant as naturalistic alternatives to traditional moral realism. The account on offer is based on a perspectival version of David Lewis's Humean "best-system" account for scientific representation. On this version, suspect concepts are vindicated (from the perspective of a given science) when their use is required for ideal systematization. Lewis's important case, his hard problem, was to provide an Humean vindication of chance insofar as it is a required aspect of ideal physics, the best way (from the perspective of fundamental physics) to give a summary description of regularities. The present essay utilizes Michael Graziano's account of consciousness as proof-of-concept. Personal-level psychology, if Graziano is correct, is best based on a brain schema evolved to systematize attention. This schema would seem to be precisely right for the perspectival best system: a "summary description" (Graziano's term also) requiring an *awareness* concept. I argue that the required concept is indeed both subjective and phenomenal so properly termed "awareness." Moreover, even if Graziano's account is not correct, the constructivist approach described illustrates a new explanatory option for physicalism about the mind.

> *It is common to see a paper on consciousness begin with an invocation of the mystery of consciousness.... [H]ere the topic is clearly the hard problem—the problem of experience. In the second half of the paper, the tone becomes more optimistic, and the author's own theory of consciousness is outlined. Upon examination, this theory turns out to be a theory of one of the more straightforward phenomena—of reportability, of introspective access, or whatever.... [T]he reader is left feeling like the victim of a bait-and-switch*. (Chalmers 1995)

> *Without MA [the brain's own description of mental attributes of awareness] a brain has no basis to detect, conclude, report, process, assign a high degree of certainty to, or have anything else to do with awareness. MA defines experiences. It defines awareness*. (Graziano 2013, 114)

Two decades after David Chalmers' early publications on the "hard problem" of consciousness, Michael Graziano's Attention Schema Theory, or AST, (2013, 2016, 2018, 2019, 2022) identifies consciousness with the brain's native model of its own attention. That is, awareness is the brain's concise and constantly updated *description* of its neural signals as they gain focus or "selective enhancement" within that brain's informational economy. The postulated representation, the attention schema, makes *introspective access* and *reportability* possible. But characterized in this way, the AST seems to make exactly the equivocation about which Chalmers warned back in 1995 (as quoted in the epigraph), a "bait-and-switch" from a promised account of experience to a theory of mere brain information access. On the contrary, I argue, if seen in a particular constructivist light this account *succeeds* as a theory of subjective, phenomenal consciousness.

First, the constructivism I recommend is the (recently very popular) perspectival and pragmatic take on David Lewis's

best-system "package-deal," his account of scientific law *together* with another nomological matter, chance. Roughly, Lewis's idea is that truth about the nomological is constituted by what ideal or "best-system" physics implies. The present constructivist proposal is a generalization to the personal psychology perspective: truth about consciousness is constituted by what the best systematization, from this personal perspective, implies. My goal is to develop a non-traditional physicalism for philosophy of mind much as metaethical constructivisms are meant as naturalistic alternatives to traditional moral realism.[1] These points will take several pages to introduce in preliminary form. Then section I details the version of the best-system account utilized here. Section II describes Graziano's postulated brain schemas, schemas that track and foster the control of attention. Importantly, according to Graziano's AST, these schemas are also the basis for the brain's native personal-perspective systematization explicitly in terms of "awareness." Section III argues that, even if not correct in physiological detail, the AST and its portrayal of personal-perspective psychology provide a compelling constructivist *model* for phenomenal consciousness. That is, *second*, the AST *succeeds* at least as a proof of concept for the constructivism I propose.

Lewis's hard problem, central to his Humean program, is to account for fundamental law together with physical probability and to do so from a basis of facts about local events, the non-modal, non-nomological, or "occurrent facts." He defines physical law in terms of the best way to systematically represent this basis given a language representing fundamental, natural properties. And this *best system*, the most efficient summary of occurrent fact, is defined to be the description ideally trading off informativeness for simplicity. Then to be a law is merely to be a generalization logically implied by that system. Furthermore, to provide significant predictive informativeness (especially about statistical regularities) while maintaining reasonable simplicity, a system may make fundamental probability claims and so claims that are not straightforwardly about the occurrent. For consistency with his Humeanism, Lewis takes these claims to count as true *in virtue of their place in the best system*. So, for example, if some version of quantum mechanics is indeed the best system for the occurrent facts and that theory implies radon decay with probability ½ in 4 days, then this fact about systematization grounds chance (probability as an objective, single-case probability) along with the laws. This by illustration is Lewis's *package deal*.

Recent versions of Lewis's best-system account relax the fundamentalist restriction: Rather than a commitment to using concepts and objective standards of physics, a *perspectival* version of the best system account is open to the *practical* conceptualizations and standards of other theoretical approaches, for example, those of biological and social science. Still, systematization remains a matter of efficient summary while just what counts as "efficient" is context dependent. Context, the "perspective" of a given theoretical approach, is supposed to determine the realm of fact to be systematized, the conceptual resources for systematization of these facts, in addition to the criteria for good system (e.g., *predictive* capacity from a theory that suitably compresses the data). Moreover, the concepts used in such a system may be modally loaded and so not straightforwardly about the occurrent. So, generalizing Lewis's package deal, claims regarding these non-occurrent posits are true in virtue of being implied by the best system—where now "best" is relativized to the perspective of a given science. Somewhat idiosyncratically, I will call the non-occurrent conceptions "constructs" of ideal theory. Hall (2009) calls them "manufactured quantities" and Dorst (2024) calls them "manufactured structures" while also pointing to Rawls.[2] Examples of such constructs may include the *free energy* of a chemical system, *adaptation* in biology, or *social harmony* in sociology, and like Lewis's *chance*, all are ultimately grounded in the occurrent given the corresponding concept's place in a best system. So, take "PBSA" to stand for perspectival, practical, predictive, and/or package-deal best-system account, *my take* on the family of recent and pragmatic variants to Lewis's original BSA. On the PBSA, truth about the nomological is, as with the BSA, constituted by what ideal or "best-system" science implies but, for the PBSA, the science need not be fundamental physics.

The attempt here is to extend the PBSA's constructivist strategy to consciousness and for a physicalist: Statements about awareness are not straightforwardly about physical matters but are true or false in virtue of what the ideally systematized personal psychology says about them. For the sake of argument, I take Graziano's postulated attention schema as a step toward such a best system. According to the AST, awareness and the personal perspective are together based on the proposed attention schema, the brain's *summary description* of attention. Roughly, the schema includes contents like "a red apple is attended," that is, "the red apple-signal is selectively enhanced." This schema is meant to track and so foster better control of attention within the brain. So, the schema systematizes attention data for the brain's use. Finally, the important point here, the AST proposes that the schema is co-opted in humans, translated for use by various cognitive functions including personal-perspective psychology reporting, for example, "I am aware of the red apple but so is my competitor in the hunt for food." The AST, then, would seem to be just what the PBSA ordered. Its hypothesized schemas purport to be salient summaries of the brain's attention data and from a perspective of great practicality for humans, one of *awareness of self and others*. Assuming the AST is correct or close enough—so that its schemas and awareness claims are indeed the basis for personal-perspective psychology and ultimately of the ideal personal-perspective systematization—the PBSA's constructivism counts the AST awareness claims as true. And true in virtue of being a part of the best systematization.

The PBSA strategy for consciousness can be motivated with another analogy (beside that of Lewis's chance and *without* his assumed basis of mere occurrent fact). Years ago, Daniel Dennett argued from behavior patterns to the truth of intentional state claims and again on a practical basis. To warm the reader up for his "*mild realism*" about the intentional, Dennett notes that visual patterns may emerge as we look at clusters of ink-jet dots on paper (Dennett 1991). Taking the patterns seriously allows an "efficient description" marking distinctions salient from the point of view of the human visual system (e.g., some patterns in the mosaic of ink marks best count as letters). Likewise, writes Dennett, thinking in terms of intentional states provides a way to "compress" the myriad behavioral data; this to the great benefit (survival value) of human social creatures. So, a native perspective, the intentional stance, has evolved because it is useful for prediction and control. Still, intentional systematization is

not indicative of fundamental ontology but only of useful posits of agent and intention. Apart from truths about the best interpretation (ideal intentional-stance systematization), says Dennett, there are no intentional facts.

This is very much the idea I would press under the PBSA banner: a construct, like chance or intentionality, may count as *mildly real* insofar as its conception is required for best systematization relative to an essential perspective (i.e., one required for practical purposes—like that of fundamental physics or the intentional stance). I am using Dennett's term, "mild realism," but it is important to be clear about what is being conflated. I take a *successful constructivism to provide a mild realism* or (better) *objectivity*: *truth defined in terms of a best system where "best" is relative to a practical perspective*. The goal here is to utilize Graziano's AST to defend an extension of Dennett's claim about the mild realism/objectivity of the intentional to a like claim about consciousness. At very least, I hope to show how the PBSA + AST serves to model *a new explanatory option*: consciousness as a construct of idealized personal psychology, viz., claims about consciousness as true or false in virtue of what the best personal psychology system says about them.

1 | Best System Accounts and the PBSA

The contemporary Humean about laws of nature (following Lewis 1973 and Earman 1984) postulates no more than mere happenings; only "occurrent" or "non-nomic" facts are supposed fundamental. If in addition the Humean attributes dispositions, laws, or chances, then (as I will explain below) that is an *interpretive overlay* to the occurrent facts and correct about the nomological insofar as it helps systematize the occurrent. Figuratively, this view assumes a patchwork or mosaic of happenings together with a best way (or best ways) to see it as patterned.

Adding chance into the theoretical mix is a big step for the Humean. Chance is something of a graded disposition so apparently non-occurrent. It is conceptually independent of actual frequency, dependent only on *chances* for frequencies, so is not straightforwardly functionalizable and/or reducible. (Note the analogy with consciousness.) Lewis's solution, the best-system account, is to package chance with laws so that chance statements are true just in case they are consequences of the axiom system best trading off (a) *simplicity* (e.g., *fewer* axioms and postulations) for (b) descriptive and statistical informativeness (i.e., *more* content).[3]

Simplicity and informativeness are language dependent. Lewis supposed that the language of ideal fundamental science uses predicates and relations tracking only natural properties. The hope is that actual physics will latch onto these privileged properties as it uncovers the fundamental reality: carving nature at its joints as the true best system emerges in the long course of scientific investigation. Still, many Humeans are as suspicious of natural properties as of "secret powers" and for similar reasons: both notions, they argue, are the result of metaphysical speculation and not needed to understand the success of science. Moreover, the supposition that there is a single best way to systematize in physics, and one humans can comprehend, may seem unwarranted. Finally, even if this skepticism is misplaced and there is a knowable fundamental physical truth, special science laws or principles will be required in practice. So, there are reasons to rework the best-system account to exclude the metaphysics of natural properties while encompassing non-fundamental science. Current best-system accounts handle these concerns by going pragmatic. Cohen and Callender's Better Best System account ("BBS") holds that best-system competitions are held on many different fields (different sciences, physical, biological, or social) and with different rules (different scientific standards) for whatever purpose is scientifically practical. Another way to go pragmatic is inspired by David Lewis's joint definition of law and chance: Barry Loewer's "package deal account" (2007, 2020) replaces natural properties with properties legitimized by their usefulness in science. The BBS and the package deal account have inspired a long list of pragmatic accounts of law. (Some recent examples include the following: Eddon and Meacham (2015), Massimi (2018), Hicks (2017, 2018), Hicks and Schaffer (2017), Bialek (2017), Wheeler (2016, 2019), Dorst (2019), Jagg and Loew (2020), Loew and Jagg (2019, 2020), Loew et al. (2023), Bhohal (2023).) The PBSA incorporates many of the practicality/salient information considerations from these accounts.[4] I set this out in 1.–4. just below.

1. The PBSA is *normally* understood to be Humean in Lewis's sense. Laws and chances are to be defined given their place in the best systematization of the basis of all occurrent fact. For present purposes, defining "consciousness" in terms of its place in ideal personal-level psychology, the basis can be relaxed to include all physical facts.

I begin with the typical buzz words used to elucidate the PBSA's notion of systematization. A best system provides a *salient summary* of the data, given in a *useful language*, and determined by the axiom system that provides the best *compromise* between *simplicity* and *informativeness/strength*. So, it *integrates* the facts for its users. Or, in Murray Gell-Mann's (1994) terms, the best system gives the best *compression* of the *salient regularities*. One may hope to spell out exactly what these *ideal science* and *practicality* concepts come to. But that is an overwhelmingly difficult matter. For present purposes only something schematic is needed. For example, take Barry Loewer's statement of what it takes for a fundamental physical theory to be best:

> [it is] true and best satisfies the criteria of simplicity, informativeness, comprehensiveness, and whatever other conditions the scientific tradition places on a final theory....
>
> (2007, 324)

This open-ended definition is unproblematic because Loewer needs only to show how some reasonable interpretation of "system" and "law" in physics may eschew the metaphysics of natural properties and governing laws.

2. Moreover, according to the PBSA, the openness of "best system" is a virtue. Different sciences may resolve what is best in different ways: the language and standards for the best-system competition depend on perspective. Here is how Cohen and Callender make the point.

> Every conceivable carving up of the world will, assuming we can make sense of simplicity, strength and balance with respect to these kinds, result in a competition, and when there is a winner, a Best System. That this should be possible in ecology is no more puzzling for BBS than that is should be possible in physics. ... [T]he basic ontological framework associated with BBS is a single world filled with events. What it adds to this ontological framework is a conception of laws as certain kinds of sophisticated summaries of the events. Now, because interests vary among summarizers, some of these summaries will be best told with vocabularies that carve up these events in different ways...
>
> (2010, 141)

Some elements of the new vocabulary may name constructs not straightforwardly definable in terms of the physical. And the resulting system might or might not include generalities so broad as to be deemed "laws" in normal parlance. The point is to have a useful summarization of the data from the perspective of the science in question.

Section II below gives the argument from Graziano's AST to the claim that personal-level psychology best "carves the world" and summarizes using the concept of *awareness*.

3. The Humean is willing to postulate only an austere basis of occurrent fact as fundamental. Again, this restriction is not required for the constructivist account of consciousness. Still, Humeanism provides a very useful analogy.

For practical purposes, says the Humean PBSA, principles of actual science are typically projected onto possible worlds. One example is the application of laws to small sub-systems of the universe (Earman 1984, 198). To study such a system, say a well isolated electron not bonded in an atom, we may think of a computationally tractable, toy universe with only one object. Of course, this lonely electron world (as it's sometimes called) surely has a best system far simpler and more informative than the best system for the actual world (plausibly detailing exactly how the one electron behaves). Still, the practical PBSA projects actual laws so that the lonely electron world can serve to model well-isolated electrons in the actual world, (Halpin 1999). Roberts (2008) makes this point for his own contextualist account of laws and in defense of Humean Supervenience.

> Lawhood is relativized not only to a possible world, but also to a context of utterance: a proposition P is a law at world w relative to context k just in case P plays the law role relative to the theory that is salient in k and that theory is true at w. ... Suppose for the sake of argument that we are living in a Newtonian universe. (I just made our context one in which the salient theory is Newton's.) Since the lonesome particle scenario is consistent with everything that can be truly called a law—in fact, it is consistent with the lawhood of exactly the Newtonian laws—it is true that if that scenario had been realized, the laws would (still) have been the Newtonian ones.
>
> (2008, 355–7)

Then, as a context sensitive account, the PBSA (like Robert's account), says that an electron would still have had negative charge, and the force laws to go with that charge, even were there no other particles and no electrical forces acting on it (evaluating the counterfactual from the salient perspective). This counterfactual stability of the electron's charge is more-or-less the definition of *intrinsic* (*pace* Langton and Lewis 1998).

Keep in mind that, on the Humean understanding, there is nothing in the lonely particle world, if viewed in a context independent way, determining that its laws are the same as those in the actual world. A possible world, on the Humean understanding, is merely a collection of events or a totality of occurrent fact. Attributing laws and chances, then, provides an interpretive overlay and hopefully one that is helpful for scientific purposes.

Anticipating Chalmers on consciousness and the needs of section III: One could interpret a world—say the actual world—as a kind of zombie world (no laws, just happenings—evaluating from a perspective emphasizing conceptual simplicity). Still, the practical way to interpret this *one world* is as having actual laws. So, for the Humean, this is a matter of *one* world having *two* interpretations. (If these were *two* possible worlds, then supervenience on the occurrent and so Humeanism would fail at the actual world.) Something similar should be said about Chalmers' zombie world, a world including all the behavior as in the actual world but without consciousness. Or so I will argue in this essay's last section: The PBSA proponent should understand the zombie world to be merely the actual world *interpreted as* lacking consciousness (an interpretation rejecting the normal, useful, and evolved phenomenal overlay). Zombies are possible, on the PBSA, only in the innocuous sense that we human organisms can coherently interpret ourselves as mere physical organisms. Relative to a perspective emphasizing only mechanical explanation, we humans are fancy biological mechanisms without further qualitative embellishment.

4. There is no need to say that all perspectives are on a par. Following Dennett on intentional systematization, the PBSA considers only *practical* perspectives. To put the point by illustration, truth from the perspective of a hypothetical demigod is of no importance. The PBSA is only interested in truth from perspectives of practical sciences, and for the present essay, interested in truth from the personal psychological perspective which is so practical for humans in everyday intercourse. Then, insofar as the concept of awareness is required for the ideal systemization from an essential perspective (again, a perspective indispensable for practical purposes), awareness may be considered an objective matter, that is, mildly real in Dennett's sense.

Recapitulating 1. to 4.: the PBSA attempts to take Lewis's empiricist law + chance package-deal as model while eliminating his fundamentalism both about natural properties and the

standards of theory evaluation. Useful scientific perspectives may invoke nomically loaded constructs required for a best systematization; truth and objectivity regarding these constructs will depend on (a) their required place in a best-system together with (b) the practicality of the given perspective. The PBSA, then, is a way of fleshing out what Dennett described years ago for what I would now call "constructed" intentional properties:

> ***[A]ll there is*** to being a true believer is being a system whose behavior is reliably predictable via the intentional strategy, and hence ***all there is*** to really and truly believing that ***p*** (for any proposition ***p***) is being an intentional system for which ***p*** occurs as a belief in the best (most predictive) interpretation.
>
> (1981, 72, emphasis in the original)

I propose to use Graziano's AST to detail how the science of consciousness may make a similar move: ***All there is*** to truly being a *conscious organism* is for that organism to be best interpreted (for the human's native predictive purposes and so relative to the human's native perspective) by a personal psychology employing a subjective and phenomenal conception of awareness.

So, my constructivist proposal requires relativization to an important/evolved human perspective: It is true that an organism is conscious (and conscious in a certain way) just in case the best systematized psychology, from this personal perspective, implies that the organism is conscious (in that way). Can there be such a best systematization? David Lewis suggested that decision theory is "a systematic exposition of the consequences of certain well-chosen platitudes about belief, desire, preference, and choice. It is the very core of our common-sense theory of persons, dissected out and elegantly systematized" (1974: 337–338). But this (partial) systematization of personal psychology is only about intentionality and not consciousness. And common-sense psychology has long been critiqued as weak and stagnant when the personal perspective and awareness are involved. For example, Paul Churchland writes this:

> Consider the rich variety of perceptual illusions, visual and otherwise. Or consider the miracle of memory, with its lightning capacity for relevant retrieval. On these and many other mental phenomena, FP [folk psychology] sheds negligible light.
>
> (1981, 73)

However, current purposes do not require that the folk do the systematizing. Instead, mammalian brains have evolved for that job. Or so says Michael Graziano: The hypothesized human brain attention schema is the basis for the self and other-regarding personal perspective together with systematization with respect to this perspective. Still, Graziano's schemas sketch and summarize *attention*, accordingly much development is necessary to see how his theory fits with the PBSA to serve as an account of *consciousness*.

2 | The Attention Schema Theory

In a series of papers and his 2013 book, Michael Graziano and co-workers develop the "attention schema theory" of conscious awareness. The theory begins with attention, that is, with signal amplification within the brain's neural system. Even simple nervous systems require competition between neural signals so that *salient* information (e.g., of a sudden noise) gain influence; this enhancement constitutes *attention* for a signal's content. Advanced creatures' nervous systems have mechanisms for top-down *control* of attention. These mechanisms allow global, task-dependent considerations to play a role (perhaps an expectation of noise, one indicating predator or prey, brings about focus on sound).

To better control attention, Graziano proposes that the brain *tracks* its own attention. "[I]n addition to doing attention, the brain also constructs a description of attention, a quick sketch of it so to speak, and awareness is that description." (Graziano 2013, 25) For the moment, set aside what it means for awareness to *be* a description. The current matter is this attention schema/brain model of attention evolved to precisely control attention (much like the body schemas hypothesized to model the body and allow fine motor control). A hypothesized native description of attention is antecedently plausible, Graziano notes, given presumptions of engineering control theory.

Additionally, in humans, the attention schema serves as a basis for introspective access to the contents of attention and so their reportability. Graziano (2016) describes this access in terms of a "cognitive/linguistic interface" which transforms the attention schema into an "awareness" interpretation broadcast throughout the brain for various cognitive matters including introspection, goal-analysis, planning, and reporting. Because an attention schema is normally known only through the latter route, most motivating examples in Graziano (and the ones I use) are of the attention schema as translated into English in terms of *awareness*. Still, because the schema is a *description*, the AST would seem to fall short of an explanation of consciousness. (Indeed, by his (2018, 2019), Graziano describes human awareness as something of an illusion of consciousness.) Some development is needed to see the AST as a full-fledged account of experience and not just of attention-as-modeled. To this end, it is helpful to use and extend a simple Graziano illustration, one that will help connect the AST to the practical constructivism of the PBSA.

Graziano has a favorite example, an apple attended by a human.

> The brain constructs descriptions of real entities in the real world. Those descriptions may not always be accurate. They may be simplified or schematized, but they generally reflect something useful to know. When the brain encodes information about the color of an apple, for example, that information relates to something physically real—wavelengths reflecting from the surface of the apple.
>
> (Graziano 2013, 21)

So, the brain represents (putative) objects and properties of the world: the apple and its color, shape, position, and so on. Such representations are typically scattered across the brain. *Attention*, on the Graziano hypothesis, binds them into an integrated whole, an *object-model*:

> There is a red apple on the branch.

So far, this is just a matter of the brain's focus on the apple and its objective properties.

Beyond the objective representations, the brain takes note of its own perspective. According to the AST, this is accomplished via a *self-model* (of the sensed physical body, matters of personal biography, goals, and personality). The final model needed for the AST is the *attention schema* posited to encode the attention relationship between self and object. Together, the models' content is rendered (by the cognitive/linguistic interface) as something like:

> I see the red apple on the nearby branch.

Or generalized to a full personal and social perspective when the schemas represent *other* humans' attention or awareness:

> *I* see the red apple on the nearby branch but my *competitor* in the hunt for food does not.

In this way, according to the AST, rudimentary first-person psychology has its roots in the attention schema (taken together with the self and object models). As Graziano argues, the attention schema is evolutionarily valuable. If a hungry human sees a red apple on a nearby branch but an even bigger one on an unreachably high branch, then it takes self-control to keep one's eyes on the accessible, if smaller, prize.[5]

Moreover, by hypothesis, the attention schema is a neat if oversimplified *summary* description of attention, while attention arises from the brain's *integrated* accounting of what is salient both within and external to the neural system. This point is especially important for purposes of the best-system account because concise description and integration help constitute good systematization. In this vein, Graziano writes that the schema is "efficient, simplified, useful, and not very accurate ..." (Graziano 2018) and "a *summary* of a deeper data set in the brain" (Graziano 2013, 81). "The attention schema theory works only so far as the proposed attention schema can be bound to other chunks of information. It is an integrated information approach to consciousness..." (Graziano 2013, 137), or finally:

> When a person says, "I am aware of X," that statement is of course a verbal *summary*. It is an *abstraction*. It is shorthand for a much richer set of information that lies deeper in the brain. Each word has so much meaning behind it that it is *like a flag that stands for an entire country*.
>
> (Graziano 2013, 109, emphasis added)

Much like chance in physics, the attention schema and its awareness conception make concise, integrated systematization possible. In Graziano's example, schemas represent the huge set of neural information regarding a person (oneself or another), an apple, and brain focus on the latter.

Still, the AST remains a minority view in cognitive science.[6] No matter. For the present proof-of-concept purpose, the goal is to show *how* the PBSA can be applied to the AST to give a *plausible* constructivism for consciousness. In this vein, two points from the last paragraphs require emphasis. (1) Graziano's theory is a close relation to other plausible accounts of consciousness and (2) it fits nicely with the PBSA.

First, the AST incorporates other important accounts of consciousness (or of the correlates of consciousness): Higher order-theory, binding, social theories, global workspace, information integration, and Bayesian/Generative Predictive Processing models. Rightly, I think, Graziano takes this as an advantage of his theory. He writes:

> In the attention schema theory, attaching the construct of awareness to a specific item—whether an apple, or a thought, or anything else—requires some method of integrating information across disparate brain areas into a single, larger, brain-spanning representation.... In this sense, the attention schema theory resembles many previous proposals in which consciousness depends on integration of information, a binding of information, a brain-wide global workspace, or a settling of networks into a single coherent state.
>
> (Webb and Graziano 2015)

Also worth mentioning is that generative/predictive processing (PP) and other Bayesian representation theories of the brain make room for the AST-like awareness representations of self and others, representations that serve efficient brain information processing. Andy Clark (2019 656–658) describes a "potentially rich contact" between PP models and Graziano's at just the point where they describe the personal psychology of awareness (e.g., the apple awareness cases above). Moreover, Clark's endorsement is based on efficient summarization.

> The upshot is that the contents that constitute our qualitative experiences are subtly responsive to our own reactive dispositions and current physiological state in ways that typically remain hidden. They function only as what Koch called "*executive summaries*," whose cognitive role is to support choice and action. In the PP story, they are elements in a *simplified* model of ourselves—a model that (like any other model learned using the PP apparatus) is under pressure to be both *accurate and concise*, capturing only the complexity needed to support *behavioral and adaptive success*.
>
> (Clark 2019, 656, *emphasis added*)

Indeed, Koch and Tsuchiya (2007) provide an early version of this PBSA-friendly practical summary idea:

> Consciousness is surmised to have substantially different functions from attention. These include summarizing all information that pertains to the current state of the organism and its environment and ensuring this compact summary is accessible to the planning areas of the brain, and also detecting anomalies and errors, decision making, language, inferring the internal state of other animals, setting long-term goals, making recursive models and rational thought.
>
> (Koch and Tsuchiya 2007, 17)

On this view, information integration—here understood as providing a *compact summary* used in practical reasoning—is essential to consciousness. Graziano has made a good case that the AST's schemas do exactly that job: summarizing the salient information about brain attention (recall that attention itself is supposed to integrate salient information) and sending the results on for further cognitive processing at the personal level of description and in terms of awareness.

This leads to my second point relating the AST to the PBSA. The schema is a representation that both defines the brain's practical personal perspective and allows efficient systematization (in summary or integrated form) from that perspective. This is exactly what the PBSA's constructivism requires. Still, the AST's personal perspective and the attention schema-as-summary provide a theory of *attention-as-modeled*. It remains unclear that this is a theory of *awareness* (properly understood) and not the bait-and-switch against which Chalmers warned. The final section of this essay addresses that skepticism.

3 | Graziano's "Awareness" as Subjective, Phenomenal Consciousness

To recapitulate, the PBSA's constructivism presupposes a useful perspective and defines truth about laws and associated constructs (like *chance*, *belief*, and, I argue here, *consciousness*) in terms of what the best system, from that perspective, implies. The constructivist proposal at the end of section I is this: An organism is conscious (in a certain way), relative to the practical and evolved human perspective, just in case the best systematized psychology, from this personal perspective, implies that the organism is conscious (in that way). Section II shows how the AST's schemas may underpin this constructivism. According to Graziano, the schemas provide a summary *of attention* but also, together with a cognitive/linguistic interface, determine the practical personal perspective and the awareness representations to efficiently describe the human organism from that perspective. "The apple appears bright red to me but my competitor for food does not yet see it" illustrates the point. Still, and this is the concern broached at the end of the last section, it remains unclear how this theory stemming from attention-as-modeled can underwrite a constructivism for *full-blooded awareness*. This section argues that the AST + PBSA does just that.

For the sake of my proof-of-concept argument, I assume that Graziano and the AST are more-or-less correct about actual common-sense personal psychology and its origin in brain schemas as interpreted by the cognitive/linguistic interface. And that actual personal psychology is normally (in healthy, properly functioning humans) more-or-less in agreement with idealized personal psychology regarding rudimentary claims. That is, if in normal circumstances the cognitive/linguistic interface interprets someone's brain schemas to include something like "I see the red apple on the branch," then this much is correct about what ideal personal psychology would say about that person in the situation.[7] Now, to make my argument, I address three crucial questions regarding the proposed constructivism. Number one, is the AST's personal perspective an essential one rather than just an accident or byproduct of evolution? (If not, its realism may be *too* mild to vindicate the significance of consciousness and avoid illusionism. That is to say: if awareness matters are to count as objective, the PBSA requires that their truth or falsity be relative to a perspective indispensable for practical purposes.) Number two, is the AST's "awareness" conception subjective and phenomenal (so truly about consciousness)? And finally, number three, can the AST + PBSA hold up against some pretty obvious objections? (E.g., even granted that the AST is conceptually subjective and phenomenal, one might object that it still leaves out *experience* and/or is unable to account for the *what-is-it-like* of consciousness.)

1. First, assuming Graziano is on the right track, the personal perspective is a required perspective for healthy humans and likely important for other species. Of course, the personal point of view is not the only useful perspective on the human or other biological organism. But personal psychology provides a native and centered viewpoint for the organism, a perspective that has evolved with mammals. According to the AST, this perspective has clear evolutionary value both for controlling attention, understanding that of others, and in humans, understanding in terms of awareness. Moreover, there is more than a contingent evolutionary justification. Given control theory, there may be something necessary about the evolved attention schema: an organism needs to track its own attention to avoid distraction. The *awareness* construct, on this view, is practically indispensable and so provides the kind of objectivity or mild realism the PBSA constructivism can provide.

Still, this last conclusion may seem overstated because the term "awareness" may be inappropriate (for the subject matter of attention-tracking) even if the AST + PBSA is otherwise unexceptional. So, the question still to be addressed: How is the attention schema's awareness concept a matter of subjective and phenomenal consciousness and not just attention-as-modeled?

2. Is the term "awareness" apt for the attention-tracking schemas postulated in the AST? A positive answer, as I understand the question, indicates that awareness is both *subjective* and *phenomenal*.

Begin with *subjectivity*. Stereotypically, this is the notion that experience is *for a subject* and that knowledge of self from the subject's *point of view* is *privileged and authoritative*. In personal psychology (intuitively and as based on the AST), awareness is explicitly *subject* oriented. The assumed AST brain schema

structure is of the form self-model + attention schema + object model, where the object of awareness is anything that can be attended, for example, Graziano's apple.

> I see a red apple to my left.

The subject of this statement, the self, is simply posited from the brain's personal-level *point of view* as the brain integrates self-regarding information (from the usual senses together with proprioception of body relative to the world and interoception of its internal state) to form self-models. The self, on the AST, is not theory or schema independent. Rather, like awareness, it is a construct of the brain's theorizing. Or alternatively, "she sees the red apple to my right." The subject of such statements, the other, is simply posited from the brain's personal-level *point of view* as the brain integrates other regarding information (from the usual senses and from various indicators most famously those about the gaze of others) to form the self-model including that of the other's attention.

Next, introspection is possible because the brain translates the attention schema for cognitive purposes. Thus, one can know that he or she sees or seems to see the apple. Such introspective access is (more or less) *privileged* according to the AST. Reading the attention schema is a matter of interpreting the brain's proprietary representation of its schemas, a matter for the cognitive/linguistic interface. For all practical purposes, it is hidden from others until one speaks or writes. Moreover, introspection is literally *authoritative*: the cognitive/linguistic interface, the interpreter of the schema, is the *author* for the first draft of personal psychology and so authors the first step toward the best system. (Of course, any mechanism like the AST's cognitive/linguistic interface is fallible, though the proof-of-concept assumption here is that it provides a good first approximation to the ideal personal perspective psychology.)

The case for *phenomenality* is more involved. Stereotypically, phenomenal experience is a matter of *appearance* that is *unified*, *ineffable*, *intrinsic*, and *irreducible*, (i.e., not identifiable with or explicable in terms of brain structure, physiological function, or biological mechanism). Awareness in personal psychology as described by the AST and as understood by the PBSA fits this stereotype well. Or so I argue in the next paragraphs.

First, according to the AST, awareness is rooted in attention-as-modeled. We have seen examples of the resulting introspective reports like.

> I see a red apple.

But other results of the translation may include (depending on the self-model):

> I imagine a red apple.

and

> I am hallucinating a red apple.

Thus understood, the awareness claims coming from the postulated brain schemas are indeed making *appearance* claims. Moreover, and second, its rootedness in attention has the by-product that experience is *unified*: According to the AST, attention integrates various modalities and across the space of the cortex to form an object model. Then attention-as-modeled, that is, awareness as described by the attention schema, inherits this integration.

Third, the *ineffability* of consciousness is a consequence of the AST: introspection leaves much of experience inexpressible. This is because the attention schema broadcast is muffled, translated in a lossy way by the AST's postulated cognitive/linguistic interface. Thus, claims Graziano, the schema loses its richness by the time it is used for introspective report. "The brain constructs a data set [the attention schema] that paints a rich picture of your own attention directed at something, the spatial and temporal dynamics of your attention, the implications, the likely effects of it on behavior. The data set is so rich and complex that it is largely unverbalizable." (Graziano 2013, 54).

Fourth, awareness is *intrinsic and internal to the experiencer*. That is, an organism's having awareness (as determined by the attention schema) is independent of matters outside the organism. Roughly, this comes to the kind of counterfactual stability described in section I, a consequence of the contextualist PBSA. A useful illustration is a lonely human, say Davidson's Swampman created spontaneously in a world where he is the only creature (so like the lonely electron). The practical thing to do is to take the actual world's perspective on these purely hypothetical models. The lonely electron has negative charge (we project actual physics principles) while the lonely human has experience depending on schema properties (assuming we project the AST). So, if Swampman happens to have a schema translating (by normal human standards) to "I am aware of a red apple on the tree," then we naturally attribute redness experience even though Swampman and Swampman's brain lack the normal human etiology.

Fifth, and most important, is the supposed *non-reducibility* of consciousness to brain states. There would seem to be a fundamental conceptual error in the reduction of a token conscious state to any brain structure, function, or mechanism. At very least, any such reductive account would seem "arbitrary" as Levine puts it when arguing for an *explanatory gap* between the phenomenal and physical.

> While appeal to the neurological structure of the state (together with an account of the overall physical structure of the relevant portions of the nervous system) explains a lot about how various stimuli cause the visual experience, and how the visual experience interacts with other states to cause both behavior and other cognitive states, the qualitative character of the experience—what it's like to have the experience—does not seem to be explained. One feels that this neurological configuration could just as easily have gone with a bluish visual experience as a reddish one; or, in fact, it could just as easily,

> for all we can tell, have gone with a state that there was nothing it was like at all to occupy. It is this contrast—this sense of arbitrariness that attends the psycho-physical reduction as opposed to the sense of intelligibility that attends other theoretical reductions—that is the core problem that goes by the name of "the explanatory gap."
>
> (Levine 2007)

Indeed, more physical or neurological detail, even complete detail, would do nothing to close the gap but only portray the human as a machine. The AST + PBSA constructivism does not attempt such a reduction; it does not identify or explain experience in terms of some brain structure, function, or mechanism. Specifically, consciousness is not identified with the attention schema. Rather, *claims* about consciousness are true or false in virtue of what an idealized or best-system personal psychology would say (something which I presuppose *is* based on the attention schemas and their cognitive/linguistic interface translation). Personal psychology's awareness language, then, is best seen as part of an *interpretive overlay* to the physical or biological facts about a human organism. Here the interpretation is a systematization from the organism's native point of view, for example, that the apple appears reddish to me. This is *not arbitrary* (the Levine concern). Rather, it is part of our evolved, efficient, but holistic understanding at the personal level of description.

However, at least in principle, this awareness interpretation is an *optional overlay* to the (assumed) totality of physical facts making up the world. Imagining a typical human as a *zombie* is just refusing to take the standard interpretive step while emphasizing a mechanical or neurological perspective. A zombie, then, is just a garden variety human interpreted as without the normal (and normally irresistible!) personal-perspective phenomenal characteristics. At issue, then, is not one normal human contrasted with a (distinct, merely possible) zombie twin but just one human organism interpreted with or without further personal-perspective qualitative embellishment.[8]

The analogy with chance on the best-system account should be clear. In contemporary physics, it is likewise indispensable to dress the occurrences up with talk about chance. But we do not identify chance with atomic configuration, mechanism, or function. Rather, at least according to the perspectival best system accounts inspired by Lewis, physics is an interpretive overlay onto the events: what counts as true about chance (relative to the physics perspective) is whatever ideal physical systematization requires. Chance and consciousness, then, are not matters in addition to the basic physical facts but are integral parts of highly useful interpretations of those facts.

This current subsection is primarily meant to show how the AST + PBSA's awareness *conception* is both subjective and phenomenal. Subsection 1's conclusion can now be restated: Insofar as the AST provides the basis for an idealized personal psychology from an essential perspective, then a mild realism about *awareness* (properly understood as subjective and phenomenal) follows.

3. The AST + PBSA's constructivism will face many objections. Some of these I see as burden of proof matters which cannot be seriously addressed in this short essay. Still, several should be noted. First, the account may seem question begging. Chalmers (1996, 86) explicitly rejects Humean accounts as problematic in the same way that physicalism is problematic. Or the account on offer may seem to suffer because explanation via unification or integration is now seen as out of vogue in the age of mechanistic explanation. Or again, the account may seem insufficiently reductionist because it depends on best interpretation while interpretation is mentalistic. Or, finally, to undermine the analogy between consciousness and chance, it may be argued that one has direct knowledge of selves and feelings while chance is (merely?) a theoretical posit. I agree that these concerns require answers. But in the current context, I only promised to show how a neglected explanatory strategy may be a serious contender in the competition. Given that all alternative accounts of consciousness have serious difficulties, this bar is pretty low. It seems to me, then, that the potential of the constructivist model on offer is strong enough that the objections just mentioned represent comparatively minor obstacles to be finessed as needed.[9]

Still, one serious objection must be addressed head-on. The concern is that the AST + PBSA model, even if *formally* or *conceptually* subjective and phenomenal, still leaves the experience out; it does not account for the *what-is-it-like* of experience. So, this model may disappoint those wanting a theory that "captures" feelings (an understanding of the kind great literature produces) or allows us to "deduce the distinctive aspects of experience" (from the AST one can only deduce *descriptions* of an experience or of an aspect of experience). Moreover, the very idea, "what it's like," is a question typically answered by reference to shared experiences. ("What is it like for an astronaut to orbit the earth?" Happily, I am told, it is like motionlessness rather than endless falling.) Still, short of brain manipulations, perhaps to induce attention schemas, there is nothing in knowing the AST that will provide such visceral understanding. But put this baldly, it is clear that "what-it's-like" questions so understood are not of the type a theoretical account is meant to answer. Similarly, we cannot expect a constructive account of justice, say in terms of what ideal impartiality requires, to provide a visceral sense for slavery's injustice. Nonetheless, the AST + PBSA account on offer is a descriptive account *of* the visceral.

To conclude, it is worth summing up this last concern and my response. The common terminology at issue, "knowing what it's like," "captures feelings," and "deduces aspects of experience" are equivocal regarding knowledge of consciousness. And this is so because there are many types of knowledge associated with the many types of neurological representation available to humans (Churchland 1985). The constructivist account on offer provides *one* such means: theoretical/discursive knowledge (and *not* visceral or empathetic understanding). For example, a person has redness-experience just in case, and in virtue of the fact that, the best systematization of personal psychology says so. If the AST is more-or-less correct, then the personal perspective and the best personal-perspective psychology are based on our native cognitive/linguistic interface interpretation of brain schemas (attention schemas together with self and object models). And so, facts about someone's redness-experience are

determined by what that person's schemas, in all their richness, are *best* interpreted to imply. Moreover, this constructivism is not tied specifically to the AST. Other theories of consciousness take the brain as self-modelling and self-interpreting (e.g., the predictive processing model briefly addressed in section II) and so may also be able to underpin the constructivist account on offer.

Conflicts of Interest

The author declares no conflicts of interest.

Endnotes

[1] I take the (contested) notion of constructivism from political and ethical theory. I understand it broadly to include ideal observer and interpretivist views as well as formal, procedural ones. The construction at issue is hypothetical. Rawls provided the model: "Kantian constructivism holds that moral objectivity is to be understood in terms of a suitably constructed social point of view that all can accept. Apart from the procedure of constructing the principles of justice, there are no moral facts." The procedure here is that of the "original position" whose idealized citizens "select first principles of justice for their society" (1980, 518–520). I propose that certain kinds of *scientific* objectivity, objectivity about scientific laws/principles and associated nomological matters (viz., physical probability, intentionality, and consciousness) are likewise grounded in a suitably constructed (idealized) scientific system. Apart from this scientific best-systematization, there are no nomic, intentional, or phenomenal facts.

[2] "Imagine that you are behind a veil of ignorance that prevents you from knowing which person you are in the mosaic, while at the same time you know the entirety of the mosaic. Using that knowledge your task is to design a set of principles—when the veil of ignorance is lifted and you figure out who you are (but also lose your knowledge of the entirety of the mosaic)—would be most predictively useful to you." (426).

[3] Statistical content is not clearly defined here. We have a rough but practical idea of how chance claims "fit" the facts from the Principal Principle identifying rational credence with chance in specific circumstances.

[4] Compare Rawls's *practicality* for the constructed principles constituting his conception of justice (as described in footnote 1): "[I]t is desirable that knowing whether these principles are satisfied, at least with reference to fundamental liberties and basic institutions, should not depend on information difficult to obtain or hard to evaluate. To incorporate these desiderata in a constructivist view, the parties [original position theoreticians] are assumed to take these considerations into account and to prefer (other things equal) principles that are easy to understand and simple to apply." Rawls explicitly trades off strength of the principles for the "practical" benefits of simplicity given the "inevitable limitations of our moral capacities and the complexity of our social circumstances" (1980, 561–562).

[5] In other words, and this is control theory in action, awareness of the bigger but inaccessible apple allows one to set it aside from consideration while picking the lower-hanging fruit. Experimental studies of spatial attention work with virtual distractors, disks on a computer screen. Some disks are "masked" so that a subject viewing the screen typically lacks awareness of them. The studies show that masked disks—the ones of which the subject is unaware—are *more* distracting and may keep the experimental subject from attending a task-relevant stimulus. The interpretation is that awareness (through the schema representation) minimizes distraction (Webb et al. 2016).

[6] It is worth noting, though, that there are neurological studies indicating where and how the attention schema may be realized in human brains. Graziano and co-workers have extensively studied the temporoparietal junction or TPJ. Damage there "can lead to neglect, a loss of awareness of objects and events on the side of space opposite to the brain lesion .... Yet the TPJ has also been implicated in attributing mind states to others" (Graziano and Webb 2017, 553).

[7] This is not to say that an ideal personal psychology can ever be perfected to remove the many limitations of folk psychology. For example, psychology from the personal perspective cannot now, and perhaps can never, explain optical illusions or the hallucinations of mental illness. Ideal personal-perspective psychology may simply have to be cognizant of its limitations. Generally, the PBSA holds that there are many perspectives of value. Human commonsense psychology is of very limited scope even while especially valuable in everyday human contexts.

[8] The zombie argument for dualism requires two *distinct* humans: One actual human who is conscious and a second possible zombie human, physically identical to the first, but without consciousness. Then the actual human would have something, viz., consciousness, *over and above* the physical character exactly shared with the zombie twin. The present constructivist proposal denies that there are two possible humans here and affirms only distinct interpretations of one and the same physical being.

[9] To begin, I would argue that the burden of proof is now on the objector's side. For example, the anti-Humean argument I cite from Chalmers is based on an old argument from Carroll (1990). But Carroll's (2018), "Becoming Humean," now sees the Roberts (2008) "lonely electron" argument as undermining his (1990) concerns with the BSA. This argument is also explicit in Halpin (2003).

References

Bhohal, H. 2023. "The Package Deal Account of Naturalness." In *Humean Laws for Human Agents*, edited by C. Loew, S. Jaag, and M. T. Hicks. Oxford University Press.

Bialek, M. 2017. "Interest Relativism in the Best System Analysis of Laws." *Synthese* 194, no. 12: 4643–4655.

Callender, C., and J. Cohen. 2010. "Special Sciences, Conspiracy and the Better Best System Account of Lawhood." *Erkenntnis* 73, no. 3: 427–447.

Carroll, J. 1990. "The Humean Tradition." *Philosophical Review* 99: 185–219.

Carroll, J. 2018. "Becoming Humean." In *Laws of Nature*, edited by W. Ott and L. Patton. Oxford University Press.

Chalmers, D. 1995. "Facing up to the Problem of Consciousness." *Journal of Consciousness Studies* 2, no. 3: 200–219.

Chalmers, D. 1996. *The Conscious Mind*. Oxford University Press.

Churchland, P. M. 1981. "Eliminative Materialism and the Propositional Attitudes." *Journal of Philosophy* 78, no. 2: 67–90.

Churchland, P. M. 1985. "Reduction, Qualia, and the Direct Introspection of Brain States." *Journal of Philosophy* 82, no. 1: 8–28.

Clark, A. 2019. "Consciousness as Generative Entanglement." *Journal of Philosophy* 116, no. 12: 645–662.

Dennett, D. 1981. "True Believers: The Intentional Strategy and Why It Works." https://philpapers.org/archive/DENTB.pdf. Originally in A. F. Heath, ed., *Scientific Explanation* (the Herbert Spencer Lectures at Oxford), Oxford University Press.

Dennett, D. 1991. "Real Patterns." *Journal of Philosophy* 88, no. 1: 27–51.

Dorst, C. 2019. "Toward a Best Predictive System Account of Laws of Nature." *British Journal for the Philosophy of Science* 70, no. 3: 877–900.

Dorst, C. 2024. "Does the Best System Need the Past Hypothesis?" *Philosophy of Science* 91: 410–429.

Earman, J. 1984. "Laws of Nature: The Empiricist Challenge." In *D.M. Armstrong*, edited by R. J. Bogdan. Reidel.

Eddon, M., and C. J. G. Meacham. 2015. "No Work for a Theory of Universals." In *A Companion to David Lewis*, edited by B. Loewer and J. Schaffer. John Wiley & Sons, Inc.

Gell-Mann, M. 1994. *The Quark and the Jaguar: Adventures in the Simple and the Complex*. Henry Holt and Company.

Graziano, M. 2013. *Consciousness and the Social Brain*. Oxford University Press.

Graziano, M. 2016. "Consciousness Engineered." *Journal of Consciousness Studies* 23, no. 11–12: 98–115.

Graziano, M. 2018. "The Attention Schema Theory of Consciousness." In *Routledge Handbook of Consciousness*, edited by R. Gennaro. Routledge.

Graziano, M. 2019. *Rethinking Consciousness: A Scientific Theory of Subjective Experience*. W.W. Norton.

Graziano, M. 2022. "A Conceptual Framework for Consciousness." *Proceedings of the National Academy of Sciences of the United States of America* 119, no. 18: e2116933119.

Graziano, M., and T. Webb. 2017. "From Sponge to Human: The Evolution of Consciousness." In *Evolution of Nervous Systems*, edited by J. Kaas, vol. 3, 547–554. Elsevier.

Hall, N. 2009. "Humean Reductionism About Laws of Nature." https://philarchive.org/rec/HALHRA.

Halpin, J. F. 1999. "Nomic Necessity and Empiricism." *Noûs* 33, no. 4: 630–643.

Halpin, J. F. 2003. "Scientific Law: A Perspectival Account." *Erkenntnis* 58: 137–168.

Hicks, M. T. 2017. "Making Fit Fit." *Philosophy of Science* 84, no. 5: 931–943.

Hicks, M. T. 2018. "Dynamic Humeanism." *British Journal for the Philosophy of Science* 69, no. 4: 983–1007.

Hicks, M. T., and J. Schaffer. 2017. "Derivative Properties in Fundamental Laws." *British Journal for the Philosophy of Science* 68, no. 2: 411–450.

Jagg, S., and C. Loew. 2020. "Making Best Systems *Best For Us*." *Synthese* 197: 2525–2550.

Koch, C., and N. Tsuchiya. 2007. "Attention and Consciousness: Two Distinct Brain Processes." *Trends in Cognitive Sciences* 11: 16–22.

Langton, R., and D. Lewis. 1998. "Defining 'Intrinsic'." *Philosophy and Phenomenological Research* 58, no. 2: 333–345.

Levine, J. 2007. "Phenomenal Concepts and the Materialist Constraint." In *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*, edited by T. Alter and S. Walter, 145. Oxford University Press.

Lewis, D. 1973. *Counterfactuals*. Blackwell.

Lewis, D. 1974. "Radical Interpretation." *Synthese* 23, no. 3: 331–344.

Loew, C., S. Jaag, and M. T. Hicks, eds. 2023. *Humean Laws for Human Agents*. Oxford University Press.

Loew, C., and S. Jagg. 2019. "Humean Laws and (Nested) Counterfactuals." *Philosophical Quarterly* 70, no. 278: 93–113.

Loew, C., and S. Jagg. 2020. "What Humeans Should Say About Tied Best Systems." *Analysis* 80, no. 2: 273–282.

Loewer, B. 2007. "Laws and Natural Properties." *Philosophical Topics* 35, no. 1/2: 313–328.

Loewer, B. 2020. "The Package Deal Account of Laws and Properties (PDA)." *Synthese* 199: 1065–1089.

Massimi, M. 2018. "A Perspectivalist Better Best System Account of Lawhood." In *Laws of Nature*, edited by W. Ott and L. Patton. Oxford University Press.

Rawls, J. 1980. "Kantian Constructivism in Moral Theory." *Journal of Philosophy* 77, no. 9: 515.

Roberts, J. 2008. *The Law-Governed Universe*. Oxford University Press.

Webb, T., and M. Graziano. 2015. "The Attention Schema Theory: A Mechanistic Account of Subjective Awareness." *Frontiers in Psychology* 23: 500.

Webb, T., H. Kean, and M. Graziano. 2016. "Effects of Awareness on the Control of Attention." *Journal of Cognitive Neuroscience* 28, no. 6: 842–851.

Wheeler, B. 2016. "Simplicity, Language-Dependency and the Best System Account of Laws." *Theoria* 31, no. 2: 189–206.

Wheeler, B. 2019. "Compressibility and the Algorithmic Theory of Laws." *Principia* 23, no. 3: 461–485.