

The Genealogy of Content or the Future of an Illusion

Alex Rosenberg¹

Received: 21 April 2014 / Revised: 1 November 2014 / Accepted: 21 May 2015 /
Published online: 12 July 2015
© Springer Science+Business Media Dordrecht 2015

Abstract Eliminativism about intentional content argues for its conclusion from the partial correctness of all three of the theses Hutto and Satne seek to combine: neo-Cartesianism is correct to this extent: if there is intentional content it must originally be mental. Neo-Behaviorism is correct to this extent: attribution of intentional content is basically a heuristic device for predicting the behavior of higher vertebrates. Neo-Pragmatism is right to this extent: the illusion of intentionality in language is the source of the illusion of intentionality in thought. Eliminativists employ the insights of all three “neo”-theses to explain why there is no such thing and why the systematic illusion that there is intentional content runs so deep.

Keywords Intentionality · Teleosemantics · Naturalization · Interpretation

Neo-Cartesians are defined, by Hutto and Satne (2014, hereafter HS) as those who hold that thought has intentional content, and that among the other things that have intentional content, including speech and inscription, it is conferred on them in virtue of their dependence on thought. The original intentionality/derived intentionality distinction that HS employ to define neo-Cartesianism is of course Searles'. There is a very good argument for the dependence of linguistic intentionality on cognitive intentionality, which they reprise: the evolutionary emergence of language and individual language learning both require having prior thoughts with content, so the latter must have content before the former. Assuming that the intentionality of language is the same as the intentionality of thought, the latter confers its original intentionality on the former by conscious and nonconscious acts of “interpretation.”

Neo-Behaviorists are initially defined by HS as those who hold that that the intentionality of thought is a matter of ascription, attribution, interpretation by others. Its existence in thought is dependent on the way thinkers are treated by other thinkers. Since this way of characterizing neo-Behaviorists subjects them to immediate and

✉ Alex Rosenberg
alexrose@duke.edu

¹ Duke University, Durham, NC, USA

obvious criticism for begging the question, it's clear that there must be something more to this view. Dennett's "Real Patterns" (1991) argument suggests that neo-Behaviorists hold that there is something *physically* distinctive about the cognition to which we ascribe intentional content, though the ascription only characterizes its afferent consequences in behavior.

Neo-Pragmatists are, according to HS, those who deny the neo-Cartesian commitment to Searle's (1983) thesis that intentional content is originally in the brain and derivatively in language and human action. They hold that intentional content starts in the culture and then infects the brain, or at least its intentionality is dependent on the existence of cultures and that cultures are not reducible to thoughts.

These theses are supposed to be exclusive, and though not logically exhaustive, at least to "cover the playing field."¹ Here is an alternative position that can be defined using these HS's three categories: Eliminativism about intentional content holds that neo-Cartesianism is correct to this extent: if there is intentional content it must originally be mental. Neo-Behaviorism is correct to this extent: attribution of intentional content is basically a heuristic device for predicting the behavior of higher vertebrates. Neo-Pragmatism is right to this extent: the illusion of intentionality in language is the source of the illusion of intentionality in thought. Eliminativists combine these three theses with the denial that there is any intentional content. In fact they employ the insights of all three "neo"-theses to explain why there is no such thing and why the systematic illusion that there is intentional content runs so deep.

One thing it's worth doing before developing the eliminativist argument, or any other argument about this matter, is to have something approaching an operational definition of intentionality, or some litmus test that all parties can agree to about what counts as the subject of the debate. Whichever characterization is chosen, will provoke dissent, but one has to start somewhere if the dialectic is not to degenerate, circle back or find its parties at cross purposes. So, let's treat, as the minimal mark of content, failure of substitutivity of co-referring terms and co-extensive predicates in the description of the intentional state. The short label for this feature of intentional states is *opacity*. There are venerable and well-known difficulties with the *opacity criterion* for intentionality. In particular some mental states are not opaque, and opacity—the failure of substitution of coreferring terms and coextensive predicates—is manifest by other states and conditions, including modal conditionals. Moreover, reference and predication are overly linguistic acts. But neo-Cartesians especially may not wish to commit themselves to the linguistic character of thought, a thesis that might beg the question in favor of neo-Pragmatists. We can try to honor this scruple by employing Fred Detske's (1995) alternatives, 'topic' and 'comment' to state our criterion for content in a way less overly linguistic. 'Topics' of thought and the 'comments' thought makes on its topics need not be nouns, verbs, adjective and adverbs, even though the vocabulary we employ to advert to them will be. Thought with a topic thereby has content, and it looks like thoughts with topics can't fail to comment on their topics at least minimally. Of course if neo-Pragmatists are right, this scruple is just a dodge.

¹ Though which playing field may seem in doubt since HS mix baseball metaphors with cricket jargon—"pitch" instead of "field," "get their innings" instead of their "times at bat," "third basers," instead of "third basemen." Shades of "silly mid-off."

Neo-Cartesianism

What psychology, and more broadly, biology, requires is, as HS note, “a story about how non-semantic causal interactions with non-contentful worldly items suffices for, or otherwise gives rise to, contentful representations.” [HS, p. 7] They immediately note that “teleosemantics is undoubtedly the most popular way” of trying to do this. Teleosemantics does so by cashing content in for biological function and cashing the latter in for variants of Darwinian adaptation.

That teleosemantics is the most popular way of naturalizing content is an understatement. Teleosemantics is the only way that neural states can be contentful cognitive states. To begin with, as a claim about the concepts we employ to identify mental states, functionalism is hardly worth arguing about. Practically every noun in natural language, most verbs, and almost every adjective, is functional, that is, has its meaning given by the causal role of what it labels, and usually its effect, not its cause. So it will be unsurprising that this will be true of ‘belief,’ ‘desire,’ and their cognates. Since one principle function of these states is, as the holism of the mental makes clear, to work together in the production of environmentally appropriate behavior, the means/ends relation between desire and belief reflects the essential teleology of thought. Of course there are desires never carried into choice and beliefs that never serve to narrow down any means to satisfy desires. But what makes thoughts into desires and beliefs is the general teleological relations among these three (mental) kinds. This is where Darwin unavoidably enters the picture. Since behavior is adaptive, either it or its causes must be the results of one or more processes of Darwinian natural selection. Darwinian processes are the only way in which adaptation, the appearance of goals, ends, purpose, etc. is created in a purely physical world.

It’s not just that intentional states have been built phylogenetically, and ontogenetically by processes of blind variation and natural selection (or better, environmental filtration). They also operate that way in real time: those that track the changing environment do so in feedback and feedforward loops that are also operated by Darwinian processes. Dennett 1995 put this point well when he distinguished four different types of creatures: Darwinian creatures whose adapted behavior is hard wired, Skinnerian creatures whose adapted behavior is learned by an ontogenetic version of Darwinian processes—the law of effect, Popperian creatures—among them the primates—whose behavior is guided by cognitive processes in which plans are randomly generated and filtered in the head, as Skinner said, enabling our ideas to die in our place. Finally, there are Gregorian creatures, who can import public signs—words—into their heads, enhancing working memory for the number and differences among random variants and providing for the construction in working memory of finer filters that select adapted behavior-plans.

Inner states produce environmentally appropriate behavior some of it as finely suited to its circumstances as speech and inscription. They do so because they are shaped by on-going Darwinian mechanisms. It is by the distinctive appropriateness of their effects in (verbal) behavior that they acquire their content retrospectively. In the same way that, much more slowly, the fetal hemoglobin gene acquired its “content”—being the gene for fetal hemoglobin.

Now here is the problem for neo-Cartesianism: the only way it can upgrade neural states to contentful states is by the process Darwin discovered. But the process of blind

variation and environmental filtration never produces behavior so specialized that it points back to one and only one *topic* and a unique *comment* as the content of the intentional states that produced it. This is Fodor's (1991) disjunction problem: there is no in principle way to distinguish the mistaken belief that there is a crow ten feet in front of me (it's actually a Currawong) from the correct belief that there is a crow or a Currawong before me (assuming I never otherwise interact with one). More broadly, the Darwinian process that shapes content can't distinguish between extensionally equivalent contents, it cannot discriminate the topics and the comments of thoughts in ways that are fine enough for the attribution of specific sentences in any natural language to mental states as their content. How much of a problem is this? Well, it depends for whom. For Searle and exponents of original intentionality, it's serious: can determinate derived intentionality—meaning, come from indeterminate original intentionality? HS, as we'll see, are prepared to settle for something like indeterminate intentionality. But it's not easy to have any "indeterminate" propositional content.

The indeterminate content neo-Cartesian is one who takes Quine's (1961) strictures seriously, while perhaps adopting some sort of computational theory of cognition. The neo-Cartesian can certainly hold that cognition involves computations over distinct tokens in thought, and that there is a distinct syntax of the tokens that these computations operate on, one that drives the processes of thought and eventually behavior. But can one hold that there is no particular proposition that any of these syntactically manipulable tokens—aka 'sentences'—express. The indeterminacy thesis cannot be the claim that intentional content is vague in the sense of being given by some indefinitely long or even infinitely long disjunction of statements, for the disjunction of all these individual statements is itself a single albeit ungainly bit of non-vague propositional content. Notice that Fodor's problem is not the impossibility of a teleosemantic distinction between the English or the mentalese sentence that there is a crow before me and the sentence that there is a crow or a Currawong before me. Content is a matter of propositions, not inscriptions in any particular language. And propositions are enough in number to describe any fact of the matter including ones that the actual sentences tokened in a language can only describe imperfectly. Propositional content wont be indeterminate; it will be indeterminate which proposition the content expresses.

Mistaking sentential content for propositional content is an even more serious mistake for the neo-Cartesian. Recall they are wedded to the Searlean thesis that the intentionality of everything else—language, action, culture, is derived intentionality, derived from original intentionality in thought. If cognitive content is given in sentences of a language, natural or mentalese, then it is not original but requires interpretation from some source that confers symbolic status on the physical matter in which the sentences are couched—in this case presumably neural spiking patterns. And now we need something else with the original intentionality to confer it on the neural patterns, so that they can convey it to public inscriptions and noises. This way lies regress of course.

HS avoid the determinate content problem by advancing a content-free kind of intentionality. According to HS the job of neo-Cartesianism is simply to give us a causal account of the origins of what they call *Ur-intentionality*. Their version of neo-Behaviorism and neo-Pragmatism will do the rest. The trouble is, *Ur-intentionality* is

just indeterminate content intentionality that was disposed of above. Here is what they tell us about Ur-intentionality:

Appeals to natural selection fail to naturalize content but they suffice to explain why certain organisms are responsive to a selective range of worldly items. Biological explanations can tell us what ancestors of a particular sort of device in fact did target and this is what fixed the range of things descendant devices now respond to, *extensionally speaking*. Thus biology provides adequate tools for making sense of something more modest than content—it provides what is needed to understand and explain responses exhibiting a kind of Ur-intentionality that results from targeted directedness of past organisms [p. 20].

What is Ur-intentionality? It can't just be selective responsiveness to stimuli. Otherwise, we'll have to credit Eric Kandel's sea slugs with intentionality. Indeed we might even have to attribute it to the Venus flytrap and other plants whose behavior is the product of "targeted directedness of past organisms." But Ur-intentionality has to be less than any state in a biological system that brings some aspect of the environment, some topic, under the thinnest possible description, the most minimal comment. In effect HS seek some halfway house between plain old biological adaptation, which is free of content, and full-blooded intentionality. To see how difficult it is to find this *tertium quid* consider that simple tropism is so to speak omniscient: the sunflower's photosynthetic apparatus tracks maximum photon exposure and there is no way to fool it into failing to do so. The point at which it becomes natural to attribute content of the simplest, most minimal kind to some organism, is when it ceases to be omniscient, when the thinnest comment it makes about its topic of attention can still be *mistaken*. Ur-intentionality is going to require the possibility of error, mistake, fallibility. It will have to be error without falsity (because the latter is semantically evaluable and so requires full-blooded, not merely Ur-intentionality). You can see the problem.

Eliminativism is the bullet-biting concession that no such *tertium quid* is available and naturalism has to get along without intentionality. But naturalism doesn't have to get along without (heuristically useful but literally false) attributions of intentionality. That is, it can help itself to a great deal of neo-Behaviorism.

Neo-Behaviorism

Long before Dennett wrote *Content and Consciousness* (1969), we were all taking the intentional stance. We just didn't know it. People take the intentional stance when the behavior they seek to explain and predict tracks, imperfectly, fallibly, non-omnisciently, a package composed of some goal state and some (immediate or mediate or even spatiotemporally distant) feature salient for the goal's realization or maintenance. It's no accident that *literal* intentional attribution to (components of) complex organisms ceases at two points: first, where behavior has a complete reductionist macromolecular explanation—molecules don't make mistakes; second, at the point in phylogeny when organism become so simple that we can identify exactly what environmental feature their behavior infallibly tracks.

As HS note, neo-Behaviorism holds to an agent/patient distinction: intentional patients are those to whom intentional content is attributed; intentional agents are those who attribute intentional content. Attribution is interpretation—it's bringing something, a topic, under a *description*, making a comment about it. Taking up the intentional stance requires that the taker have rather rich intentional states itself. There is a circle here, one tolerable to eliminativists, of course. We hold that the subjects of our attribution lack intentional content (though it's convenient for managing our relations with them to attribute it to them), and agents—we, who attribute intentional content, lack it as well, and mistake what goes on in consciousness² for our own intentionality. One way to break the circle is by invoking a well-grounded Ur-intentionality, one that starts simple and can be upgraded all the way to intentional stance-taking. Neo-Behaviorists won't need to wait for the full upgrade in order to defend their views against the complaint that they are changing the subject from explaining content to explaining why its convenient to believe in content. This is how neo-Cartesianism and neo-Behaviorism are supposed to work together. But, without Ur-intentionality, the teamwork can't get started, and as suggested above, no such Ur-intentionality is in the offing.

It seems to me that HS in fact accept the conclusion that Ur-intentionality won't work. Recall their description of it as “result[ing] from targeted directness of past organisms.” Yet their complaint about neo-Behaviorists is that “the resources they call on—while they do not assume the existence of content—are nevertheless much too thin to explain how contention and content-ascribing capacities could have come on the scene.” Why? Because neo-behaviorists (their example is Quine) “restrict [themselves] to a limited set of tools.” These tools are limited to “invoking evolutionary consideration in order to explain why we react in perceptually similar ways to salient stimuli.” Isn't this inadequate tool kit the same one HS employ to try to build Ur-intentionality?

The eliminativist acknowledges the neo-Behaviorist's observation that intentionality is a predictively useful tool, a stance we employ. It's not just useful, it's indispensable for creatures like us. The eliminativist admits this too, but nevertheless denies that the attributions neo-Behaviorism diagnoses are true. The stance is indispensable, but the facts are otherwise. The indispensability of the intentional stance is an important insight brought to us by two generations of work shared out among evolutionary anthropologists and evolutionary psychologists, game theories and experimental economists. It's worth sketching to show why a stance can be indispensable even when its attributions are false. Another reason this is worthwhile is to see what the eliminativist salvages from neo-Pragmatism.

The attribution of intentional states to others is part of the solution to a severe design problem that faced *Hominins* when they emerged on to the African savanna. Alternatively it is the breakthrough that enabled them to move from the bottom of the food chain on the savanna to the top. Either description—design problem solution or opportunity-creating break through—is equally apt. The rain forests had receded and/or the other primates had pushed *Hominins* out to their fringes on the savanna. What *Hominins* required was the ability to compete against megafauna, initially to secure

² Which after all is just a matter of more tokens succeeding each other, this time phenomenal ones, some in a syntactic order, others not, but none having any intrinsic *aboutness* about them.

scavenging opportunities, later to bring down the fauna themselves. What provided it was coordination and cooperation.

They already had an ability to predict some behavior of conspecifics and other vertebrates, an ability we now label ‘theory of mind,’ though we cannot take this label too seriously without attributing second order intentionality—thoughts about others’ thoughts about things—to elephants, dogs, cetaceans, tamarinds, and several other primates. Basically, the theory-of-mind ability is the disposition to treat another organism as a goal directed system with very rough hypotheses about the environmental factors to which its behavior is sensitive and about immediate end-state towards which its behavior is directed. If there are topics and comments in the brain states that predict behavior need not be *opaque*, and so don’t drag in content.

Once *Hominims* hit upon conventions that exploited their theory-of-mind ability together with naturally occurring gestures, grunts and other signs of their own situations, group strategies emerged that could win out against megafauna. The conventions soon become established as language of course, and a virtuous cycle emerges in which more effective, more efficient hunting provides more protein, which build bigger brains, while reducing digestive gut investment, enables *Hominims* to imitate one another better, learn, make use of the very long period of childhood dependence (hitherto a huge fitness-cost) for education, thus fostering the ability to predict and coordinate behavior, first in pidgins, and finally full blooded language—with the attribution of explicit propositional content.

The Darwinian genealogy of language makes it evident what the intentional stance does for us, and why it is indispensable. It also explains why it is much like other adaptations—a quick and dirty solution to a design problem (or opportunity), one that capitalizes on pre-existing dispositions, that gets entrenched so that it becomes hard or even impossible later to give it up even when it becomes less adaptive, or when an alternative strategy that does the same job better, becomes available. Of course nothing does better than folk-psychology yet, or at least nothing is more adaptive for getting along in human society, at this point. But, first of all, the glass is half-empty: folk-psychology with its indispensable attributions of content—doesn’t do any better than it did in the time of Homer, and it doesn’t do well enough to protect us from crazy people or help them recover; second, in some special areas content-free neuroscience is already better at explaining and predicting behavior; third, advances in neuroscience will eventually reveal the wrongness of intentional content attribution, as science has revealed the wrongness of color-attributions, folk physics, and other adaptively useful fictions.

Neo-Pragmatism

The teleosemantic version of neo-Cartesianism is right to this extent, if there is intentional content, it can only be owing to the operation of a Darwinian process. Neo-Behaviorism is right to this extent: attributions of intentional content are made in order to explain and predict the behavior of others. The extent to which these two doctrines are right does not include that there are such brain states bearing intentional content. The highly adapted character of the neural states that control behavior, and the great heuristic value of the assumption that such states have intentional content, doesn’t

entail that there are any contentful states. The intentional stance needs intentional agents and intentional patients to be armed only with Ur-intentionality. That's what neo-Cartesianism gives them. How the agents end up fooling themselves into attributing more is the job of neo-Pragmatism. To a large extent the contribution of neo-pragmatism to the story is foretold in the Darwinian genealogy of language as part and parcel of the evolution of human culture.

HS write,

According to third basers [i.e. third basemen, neo-Pragmatists], we can only make sense of contentful thinking in the context of shared ways of life in which social norm compliance is developed, maintained and stabilized through practices. Such practices are not only based on our shared biology but in social engagements and cultural devices that evolved over time, especially linguistic tokens, the *primary* bearers of semantic content [p. 13, emphasis added].

That is correct. But naturalism tells us that the social norms, like all norms, will be instrumental, hypothetical imperatives, grounded in means/ends regularities that emerge randomly and get fixed into conventions, ones not much different from 'drive on the left or the right, but not both.' We have David Lewis (1969 *Convention*, Harvard, 1969), and Skyrms 2010 to thank for an introduction to how these norms get developed, maintained and stabilized, as solutions to Darwinian design problems. And the quite natural, unmysterious scenario of how gestures and grunts—signs—get shaped into commands, instructions, requests, agreements, and the words, phrases and sentences in which they are expressed, requires no intentional content whatsoever. So, HS are mistaken when they conclude that "no one doubts that a capacity for social conformism will form at least part of the best explanation of how human [contentful] cognition did...come into being [p. 14]. Too strong: a capacity for social conformism forms part of the explanation of how the illusion that cognition has content came into being, But as for cognition itself, well, HS are clear enough that it has to be prior to the capacity for social conformism, because it's required in order to socially conform.

Neo-Pragmatism is right to this extent: semantic evaluability (truth/falsity/quantification/meaning/reference/subject/predicate, etc.) is dependent on language, speech and writing, inscription and sounds. All these are social and cultural artifacts, ones that emerged as language evolved from grunts and gestures. Semantic evaluability is created by social norms. Once they are up and running, or better, at the same time as they get up and begin running, semantic evaluability begins to be attributed to cognitive states. Why does it? Neo-Behaviorism has one answer: it's adaptive, it has survival value. Neo-Cartesianism has another answer: intentionality must originate with the mental causes of speech and writing, noises and inscriptions. But if neo-Pragmatism is right, and intentionality can only originate in the effects of cognition, in behavior, then together neo-Cartesianism and neo-Behaviorism must conclude that intentional content is just a useful fiction. And neo-Pragmatism explains how social and cultural factors foisted the fiction on us.

Eliminativism and Interpretation

Intentionality is a matter of interpretation: for a thought or a sentence to be about something is to bring what it is about “under a description.” That much is clear and uncontroversial when it comes to derived intentionality. There is nothing intrinsic to the red octagons we see at intersections that scream out at us “Stop!” They do so because of our interpretation, because we classify, categorize, identify them as, bring them under the description of ‘stop sign.’ It’s because descriptions are linguistic and language is social that neo-Pragmatism seems able to reach into the brain and plant intentionality there as children begin to articulate their thoughts to others. It’s only by interpreting behavior, bringing it under descriptions that identify its goal and the environmental conditions that make the behavior appropriate to its goal, that we can predict and control one another’s actions. That’s the lesson of neo-Behaviorism. Naturalism, here known as neo-Cartesianism, requires that interpretation end at the brain states, the neural circuitry, that carries original intentionality, that does the interpretation of behavior—actions, soliloquies, inscriptions, paintings, compositions, performances, etc. And this is why eliminativism is compelling.

Original intentionality cannot be a matter of one last interpretation. Neural circuits don’t get interoperated, can’t get interpreted, on pain of regress. But without interpretation they can’t have the sentential content they need to confer it on everything else that does come under a description, everything that has derived intentionality. There are any number of reasons why neural circuitry can’t get its intentionality by interpretation. The most obvious is the regress such a process would produce: why does the neural circuit, that ‘Paris is the capital of France’ carry that content, why is it about Paris? Call the circuit N’. Suppose N’s aboutness is a matter of interpretation: some further process of neural consolidation stored that information in circuit N’ and now another neural circuit, N” reads N’ as containing that information. How? By interpreting the synaptic connections in N’ or the spiking patterns between them? This hypothesis doubles our trouble: now we are committed to another neural circuit N” with double the interpretative duties: N”’s content includes being *about* circuit N’—bringing it under the description ‘being about Paris’ and N”’s content also has to also include being about Paris. Without the double aboutness how could the interpreting device in the brain, N”, interpret N’ as being about anything in particular. This is a *Homuncularism* problem that teleosemantics sets out to solve by forgoing N”, but cannot. Its near-miss is significant. When it comes to human action teleosemantics can give us all the predictive/explanatory power that neo-Behaviorism could hope for. It just can’t give though the sentential content that neo-Pragmatism insists rightly is required for intentional content.

Original intentionality cannot be a matter of interpretation of one part of your brain by another. Still less could it be what paleo-Cartesianism suggests—a matter of interpretation of one part of your brain by an immaterial soul, self or mind.

Some philosophers, following Searle (1983, e.g. Horgan and Tienson 2002), attribute the original intentionality to consciousness, to the introspectively accessible qualitative aspect of experience: we just have the thought, the feeling, the intuition, the conceptually unmediated knowledge that our thoughts are about things. It’s easy to accept this claim, every fiber of our conscious being is demanding we do so. But

it is hard to see how it could even be helpful, let alone correct. After all, conscious ideas, images, silent speech, etc are just more tokens in the brain, and as such bear all the same problems that the rest of the brain has in bearing the burden of unexplained original intentionality.

Original intentionality can't be somewhere else—in another person, mind, brain—interpreting the relevant part of your brain. That's not original. Another way to think about the problem is provided by Watson, IBM's *Jeopardy*-winning computer. Watson wins at *Jeopardy* without anything in its CPU or data storage having any original intentionality. The intentionality of its output is "derived," as is the intentionality of its inputs. And there is no intentionality in the intermediaries between Watson's inputs and its outputs. Watson doesn't need any to do the job. But what states with original intentionality are conferring the derived intentionality on Watson's internal states? The programmers' brain states, their neural circuitry? How did they acquire it? Why do they need any either?

The eliminativist insists that there is a clear solution to all these problems raised by the notion that one clump of matter—the brain—could have original intentionality about another clump of matter—what we are thinking *about*, say Paris or the moon or Alfa Centuri. There is no original intentionality. The flat denial of original intentionality (and with it of course, derived intentionality, all intentionality) is the conclusion to be drawn from all the unsuccessful programs of research, among neo-Cartesians, neo-Behaviorists, and neo-Pragmatists, to provide a non-circular, let alone a naturalistic account of content. Eliminativism stands on the shoulders of all those philosophers who have conscientiously sought to explain intentionality's possibility, whether as team players, the way HS advocate, or as soloists.

Eliminativism is however, left with one very large problem. HS recognize it, though they lay it at the feet of neo-Pragmatists. It is in fact a problem for any naturalistic account of content, including ones that explain it away. They write,

Third basers [neo-Pragmatists] must always bear in mind that truth-telling practices are special in as far as they involve not only socially responding to things but doing so by representing them as being this and so, independently of what we say about them.

Truth is not a matter of convention, construction or consensus. It's not epistemic, and pragmatism's notion that truth is what is agreed to at the end of enquiry doesn't cut it for HS (nor for any realist about content, *pace* Rorty 1979).

In contrast with other intelligent dealings with the environment, these content-involving practices [calling a belief true or false] contain a special sense of going wrong: this is not just what is acceptable for a community but being correct or incorrect according to how things are anyway. These practices differ essentially from ways of dealing with the world that do not represent it [p. 26].

Where could access to this kind of truth come from? Not from the propositions that sentences express. Naturalism cannot help itself to causal contact with abstract objects, and that's the only kind of contact there is. Truth is going to have to be a property of sentences. Its only source is a mental state with original intentionality, one that

compares the content of thought to the world and finds a correspondence, or a truth-maker for its content, and then extends the semantic property from its thought to the noises and inscriptions those thoughts bring about in speech and writing. But HS cannot help themselves to this sort of original intentionality, and neither can any naturalistic account of the genealogy of content.

So long as philosophers hope to find a realist theory of content—something that will reduce semantic evaluability to natural facts, they can hope that the theory will enable them to cash in truth. But once we turn our backs on intentionality, eliminativists will need to find a way of avoiding the critique that we can't express the truth of our theory, or even that it has meaning. I think we can do it, but it will require another few pages.

References

- Dennett, D.C. (1969). *Content and Consciousness*. Routledge.
- Dennett, D. (1991). Real patterns. *Journal of Philosophy*, 88, 27–51.
- Dennett, D. (1995). *Darwin's dangerous idea*. New York: Simon and Schuster.
- Dretske, F. (1995). *Naturalizing the mind*. Cambridge: The MIT Press.
- Fodor, J. (1991). *The theory of content*. Cambridge: MIT Press.
- Horgan, T., & Tienson, J. (2002). The intentionality of phenomenology and the phenomenology of intentionality. In D. Chalmers(ed.), *Philosophy of Mind: Classical and Contemporary Readings*, Oxford University Press.
- Hutto, D., & Satne, G. (2014). The natural origins of content. *Philosophia*, present issue.
- Lewis, D. (1969). *Convention*, Harvard.
- Quine, W. V. O. (1961). *Word and object*. Cambridge: MIT Press.
- Rorty, R. (1979). *Philosophy and the Mirror of Nature*, Princeton University Press.
- Searle, J. (1983). *Intentionality: An Essay in the Philosophy of Mind*, Cambridge University Press.
- Skyrms, B. (2010). *Signals, evolution, leaning and information*. Oxford: Oxford University Press.