THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Semantic inferentialism and the evolutionary argument against naturalism

Citation for published version:

Collin, J 2013, 'Semantic inferentialism and the evolutionary argument against naturalism', *Philosophy Compass*, vol. 8, no. 9, pp. 846–856. <https://doi.org/10.1111/phc3.12062>

Digital Object Identifier (DOI):

[10.1111/phc3.12062](https://doi.org/10.1111/phc3.12062)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Philosophy Compass

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Semantic Inferentialism and the Evolutionary Argument Against Naturalism

1 The Evolutionary Argument Against Naturalism

Alvin Plantinga's evolutionary argument against naturalism (EAAN)¹ seeks to establish that the naturalist who accepts evolutionary theory is committed to a pernicious and encompassing scepticism regarding the reliability of her cognitive faculties. The argument is distinctly epistemological in character; Plantinga's goal is not to show that there is some metaphysical incompatibility between naturalism and evolution, but rather, as Plantinga sometimes puts it, that the conjunction of naturalism and evolution cannot be *rationally believed*.

'Naturalism', as Plantinga understands that term, is the claim that 'there is no such person as God, or anything at all like him' [Plantinga 2008: 19]. Naturalists, needless to say, reject the notion that God, or some other non-natural person, has so orchestrated the course of evolution so as to culminate in the existence of creatures made in his image as persons capable of knowledge. As unattractive as the invocation of divine providence will be to naturalists, Plantinga's claim is that something of this sort is necessary if we are to avoid committing ourselves to scepticism. Without this kind of guidance, so the thought goes, there is no reason that creatures who form beliefs will be so constituted as to form *true* beliefs, and hence in all probability they will not reliably form true beliefs. I will present the EAAN and provide a survey of a theory of meaning known as *semantic inferentialism*. I will argue that if one accepts semantic inferentialism then Plantinga's motivation for the EAAN is undermined.

Why think that without invoking divine guidance we are unjustified in taking our beliefs to be reliable? Plantinga begins by considering materialist naturalism. On materialism beliefs either *are* or supervene upon neuro-physiological properties or events (henceforth 'NP properties' or 'NP events'). The first claim is one of identity between NP properties or events and beliefs; beliefs *reduce* to NP properties or events. The latter is a nonreductive form of materialism in which there are two kinds of properties, mental and non-mental, but mental properties are *determined* by or strongly supervene upon non-mental ones. If beliefs are or supervene upon NP properties and events, then those NP properties and events that constitute beliefs will have *contents*.² A content C is determined by an NP property or event, or, more likely, some Boolean combination of NP properties or events, for instance:

$$(P_1 \ \& P_5 \ \& P_{98} \ \& \dots) \vee (P_7 \ \& P_{12} \ \& P_{33} \ \&) \vee (P_2 \ \& P_3 \ \& P_{4^{27}} \ \&) \vee \dots$$

where P_i are NP properties. At some critical level of complexity, NP structures come to constitute beliefs or the supervenience base for beliefs, and, hence, at this critical level of complexity NP structures come to determine *contents*. On materialism then NP structures are responsible for two distinct things: on the one hand they determine

mental contents, on the other hand they cause behaviours. Given processes of natural selection, NP structures that cause adaptive behaviour are selected for, and it is these NP structures that determine contents in extant creatures, but, crucially, *what* contents these structures determine is simply irrelevant to their adaptiveness. As Patricia Churchland has famously noted,

Boiled down to essentials, a nervous system enable the organism to succeed in the four Fs: feeding, fleeing, fighting and reproducing. The principle chore of nervous systems is to get the body parts where they should be in order that the organism may survive. ... Truth, whatever that is, definitely takes the hindmost. [Churchland 1987: 548-9]

One can think of this in terms of the existence of a function f : NP $\mapsto$ C from Boolean combinations of sufficiently complex NP properties to contents. In everyday thinking one usually supposes that this function will preserve the relations we ordinarily take to obtain between the content of a belief and its role in behaviour. So, for instance, beliefs whose content is about the location of food, combined with a desire to eat will ordinarily issue in appropriate actions such as moving to the location of the food. However, given materialist-naturalism, there is no *a priori* guarantee that these relations will obtain. Moreover, evolution provides no reason to alter this assessment, for NP structures are selected for because of their causal efficacy, not because of the contents they determine. True beliefs are not selected for, rather, causal roles are, and there is no reason on this picture to think that the connection between causal role and belief content is anything but arbitrary:

Granted, the structure in question helps cause adaptive behavior. But that doesn't so much as slyly suggest that the content that gets associated with the structure is true. As far as its causing the right kind of behavior is concerned, it simply doesn't matter whether the content, that associated proposition, is true or false. At this point, as far as the truth or falsehood of the content that arises, natural selection just has to take potluck. (Not that it minds—it's interested, so to speak, just in adaptive behavior, not in true belief.) ... Given naturalism, it would be sheer coincidence, an enormous cosmic serendipity, if the content that is associated with adaptively useful NP properties should also turn out to be all or mostly true content. [Plantinga 2008: 40]

As there is no more reason to think any given belief true rather than false, Plantinga suggests taking the probability of any given belief being true as being .5. Given this, the probability that, for 1000 *independent* beliefs, at least three quarters of them were true would be less than 10^{-58} . At any reasonable measure of reliability, such a creature could not be said to have reliable belief-forming faculties, hence the probability that one's cognitive faculties are reliable given naturalism and evolution is low. In symbols:

$$P(R/N\&E) \text{ is low}$$

where R is the proposition that our belief-forming faculties are reliable, N the proposition that naturalism is true, and E the proposition that evolutionary theory is true.

The second step in the EAAN is the claim that a person who accepts naturalism and evolutionary theory and comes to see that $P(R/N\&E)$ is low, acquires, by means of that recognition, a *defeater* for R—*viz.* she comes to acquire a belief that, by virtue of holding it, precludes rationally believing R. By analogy, Plantinga imagines

a drug XX that is known to render cognitive faculties unreliable in 95 percent of those who ingest it. For these 95 percent, ingesting XX results in their believing more false propositions than true. Say that you come to believe that you have ingested the drug; it seems clear that you have a defeater for the belief that your cognitive faculties are reliable.

Three points are in order here. Firstly, because R involves *all* of one's cognitive faculties, one cannot appeal to the deliverances of those faculties, or some subset of one's faculties, to deflect the defeater, for this involves relying on the very faculties that are impugned. Secondly, and relatedly, anyone who has a defeater for R thereby has a defeater for any belief that is a deliverance of R, which is to say *all* of one's beliefs. This includes the conjunction of naturalism and evolution. Finally, because commitment to N&E entails commitment to a defeater for N&E, N&E is self-defeating.

Plantinga decomposes the EAAN into four steps:

- (1) $P(R/N\&E)$ is low.
- (2) Anyone who accepts N&E and comes to see that $P(R/N\&E)$ is low has a defeater for R.
- (3) Anyone who has a defeater for R has a defeater for all her beliefs (including N&E).
- (4) Since commitment to N&E entails commitment to a defeater to N&E, N&E is self-defeating and cannot be rationally believed.

(1) is a consequence of Plantinga's claim that we should estimate the probability of any given belief being true (given naturalism and evolution) as being .5. We will call this claim (1*):

- (1*) Given naturalism and evolution, it is no more probable that, for any adaptive belief, it will be true rather than false.

This in turn is motivated by a further claim:

[S]uppose materialism were true: then [i]t is by virtue of the NP properties, however, not the content, that the belief causes what it does cause. ... It isn't by virtue of the content of this belief; the content of the belief is irrelevant to the causal power of the belief with respect to behaviour. [Plantinga 2012: 268].

We will call this claim (1**):

- (1**) Given naturalism and evolution, the probability that there is a non-arbitrary connection between the content of a belief and its role, is low.

It is the contention of this paper that semantic inferentialism, if true, would undermine (1**). In what follows we will see why this is so, and assess its bearing on the EAAN.

2 Responding to the EAAN

There are, broadly speaking, two ways in which the EAAN has been challenged. The first is to deny that knowing that $P(R/N\&E)$ is low constitutes a defeater for naturalism. The second is to deny premiss (1), that $P(R/N\&E)$ is low. It is this second option that we shall explore here. Any cogent response to (1) must then involve a

theory of content in which there *is* some principled connection between the behaviours caused by an NP structure and the content of the belief that structure determines. This alone will hardly be sufficient however, for the function from behaviours to contents may be such that, although principled, associates mainly false beliefs with adaptive behaviours. In addition to linking behaviour to belief content then, the theory must do so in a way that links specifically *adaptive* behaviour to *true* contents.

Robustly empiricist theories of mental content, such as Dretske’s *indicator semantics*³ and Millikan’s *teleosemantics*⁴ are one such attempt to forge a link between the adaptiveness of a behaviour and its content. Plantinga [2011] discusses these elsewhere, but they are not the only semantic theories that seek to establish a principled connection between the content of a belief and its role in action. *Semantic inferentialism* (SI) is another such theory, one that poses a more difficult challenge to the EAAN. SI, as it is developed by its most prominent expositor Robert Brandom, involves an incredibly wide-ranging and multifaceted network of claims. Luckily, for our purposes, it will suffice to survey its most general features.

2.1 Semantic Inferentialism

The EAAN pertains to the content of beliefs, but what makes a belief a belief and what makes a belief have a particular content? A trained parrot’s reliable disposition to respond differentially to red stimuli by squawking “Red!” is not sufficient for it to have a true belief about the colour of the objects it is reliably responding to nor indeed for it to have a belief at all—even if it superficially resembles believing creatures in some respects. What missing ingredient then is required to elevate the creature that reliably differentially responds to various stimuli to a creature that has *beliefs regarding* the stimuli to which it responds? One answer begins by noting that the parrot lacks *comprehension*: it has no mastery, no understanding of redness; it lacks the concept RED. This is alright as far as it goes, but we have so far only succeeded in giving a series of labels to the parrot’s deficiencies. If we are to understand what it is lacking, we must somehow unpack the consanguineous notions of *comprehension*, *understanding* and *concept possession*. SI is, at its core, an attempt to do just this. Robert Brandom provides the canonical statement of the view:

To grasp or understand a concept is ... to have practical mastery over the *inferences* it is involved in—to know, in the practical sense of being able to distinguish, what follows from the applicability of a concept and what it follows from. The parrot does not treat “That’s red” as incompatible with “That’s green,” nor as following from “That’s scarlet” and entailing “That’s colored.” Insofar as the repeatable response is not, for the parrot, caught up in practical proprieties of inference and justification, and so of the making of further judgements, it is not a conceptual or cognitive matter at all. What the parrot and measuring instrument lack is an appreciation of the significance their response has as a reason for making further claims and acquiring further beliefs, its role in justifying some further attitudes and performances and ruling out others. Concepts are essentially inferentially articulated. Grasping them in practice is knowing one’s way around the proprieties of inference and incompatibility they are caught up in. What makes a classification deserve to be called *conceptual* classification is its *inferential* role. [Brandom 1994: 89]

Whereas broadly Tarskian approaches to semantics give explanatory privilege to *truth conditions*, take comprehension to consist in ‘grasping truth conditions’ and go

on to explain inference in terms of truth conditions, SI begins with a basic notion of inference and uses it to do its explanatory work. According to the semantic inferentialist picture it is impossible for a creature to possess a lone belief. Grasping a concept requires a practical mastery of its inferential relations to *other* concepts—what it follows from, what follows from it, what it precludes and what precludes it. To grasp one concept, a creature must grasp *many*. As Brandom puts it, ‘Cognitively, the grasp of one concept is the sound of one hand clapping’ [Brandom 2000: 49].

Brandom, following Wilfred Sellars’ usage, calls the kinds of inferences that make up the content of a concept *material* inferences. Material inferences are those inferences that are licensed by virtue of the *content* of the concepts involved in the inference. From “Glasgow is west of Edinburgh” one can infer “Edinburgh is east of Glasgow”, from “The paint is red (and monochromatic)” that “The paint is not green”, or from “It is raining outside” that “The ground will be wet”. Grasping a concept involves a practical mastery of the material inferences it is implicated in, an ability to *use* the concept appropriately in the practice of giving and asking for reasons and discerning one’s commitments. Inferences of this sort are not logical inferences, which is to say that the practice of making inferences is not limited to those beings who have some logical mastery. Restricting the domain of inference-making to logical discourses involves treating the above inferences as *enthymematic*; invalid until supplemented with appropriate conditionals. “Edinburgh is east of Glasgow” can only be validly inferred from “Glasgow is west of Edinburgh” with the additional premiss “For all x, y [if x is east of y then y is west of x]. Similarly, the conditionals “For all x [if x is red (and monochromatic) then x is not green]” and “If it is raining then the ground will be wet” are required to make the second and third inferences formally valid respectively. For the semantic inferentialist, material inference is fundamental, logical vocabulary (paradigmatically the conditional) is a means of *making explicit* the inferential relations that we endorse in practice (and, if one should so desire, inferential relations that we do not endorse). Explicit logical competence is not a prerequisite for inferential competence; there can be discursive practices that are not logical practices, but logical vocabulary is a means by which we can codify our discursive practices. By making explicit the inferences we could otherwise only adhere to in practice, logical vocabulary allows us to lay bare our inferential commitments, bringing them under rational scrutiny in a way that would not be otherwise possible. Material inferences are non-monotonic; supplementary or collateral premisses can turn good material inferences into bad ones. ‘It is raining outside, therefore the ground will be wet’ is a good inference, but ‘It is raining outside *and* the ground is covered with tarpaulin, therefore the ground will be wet’ is a bad one. ‘It is raining outside *and* the ground is covered with tarpaulin *and* the tarpaulin has been perforated, therefore the ground will be wet’ is, on the other hand, a good inference.

The content of a belief is constituted by the material relations it stands in to other beliefs. These relations can be classified into three fundamental kinds: *commitment-preservation*; *entitlement-preservation* and *incompatibility*. A proposition p preserves commitment to a proposition q , if commitment to p requires commitment to q . The proposition ‘That is a vixen’ preserves commitment to ‘That is a fox’, for anyone who is committed to the first is thereby committed to the second. A proposition p preserves entitlement to a proposition q if commitment and entitlement to p provides *prima facie* entitlement to q . A person committed and entitled to the proposition ‘It is raining outside’ is thereby *prima facie* entitled to the proposition ‘The ground will be wet’. Finally, a proposition p is incompatible with a proposition q just in case, if a person is entitled to p then she is not entitled to q . A commitment to ‘The paint is red (and monochromatic)’ precludes commitment to ‘The paint is green’.

Importantly for our purposes, semantic inferentialism involves a broad conception of what can play a role in material inference. The ‘inferential’ connections relevant to the content of a concept involve not only connections between separate propositions but also between their *circumstances* and *consequences* of application:

[T]he broad conception includes the possibility of noninferential circumstances and consequences of application. In this way ... the specifically empirical conceptual content that concepts exhibit in virtue of their connection to language entries in perception and the specifically practical conceptual content that concepts exhibit in virtue of their connection to language exists in action are incorporated into the inferentialist picture. Conceiving such inferences broadly means conceiving them as involving those circumstances and consequences, as well as the connection between them. [Brandom 1994: 131]

Inferential connections in this broad sense include not just connections between propositions but also afferent connections—the input of stimuli, reliable differential responses to kinds of objects—and efferent connections—the agent’s actions. Part of what constitutes our concept of redness is red perceptions (the other part is how these red perceptions are integrated into a network of relations of commitment-preservation, entitlement-preservation and incompatibility). Brandom calls the view that this broad inferential articulation is sufficient to determine conceptual content ‘strong inferentialism’. This is the view developed and defended by Brandom. ‘Hyperinferentialism’ holds that the inferential connexions that obtain between propositions are sufficient to determine content. Finally, ‘weak inferentialism’ is the view that inferential relations are a *necessary* constituent of conceptual content, but not sufficient. On weak inferentialism, if two contents were involved in different inferential connections, including behaviours, they would have different meanings in virtue of this fact.

SI’s applicability to the EAAN is probably becoming clear. Recall that it is via (1**) that (1*) is motivated and, hence, (1) itself is justified. (1**) is supposed to be plausible on the naturalist view because the role of a belief and its content can come apart. On SI however, the content of a belief is, *perforce*, relevant to its content; for what a belief *does* and what a belief *is* are two sides of the same semantic coin. It is only in virtue of possessing the right sort of content that a belief could contribute to fitness. On SI the justification for (1**) is undercut.⁵ The EAAN gains its cogency from the claim that, according to naturalism, there is no principled function from the behavioural role of belief states to their contents, or that, even if there is, one has no means of specifying this function or supposing that it associates adaptive behaviours with true beliefs. SI however involves the claim that there is a principled way in which contents are associated with behaviours, and specifies what this association is. This complicates the picture for the defender of the EAAN; she can no longer suppose *from the outset* that, given naturalism, there is no more reason to hold any behaviourally felicitous belief to be true than false, for SI holds that to specify the inferential role of a belief *just is* to articulate its content.

3 Responses

3.1 Incorrect?

Of course, the simplest response for the defender of the EAAN is to deny SI. It is, after all, a minority view; most of the theorists who earn a living thinking hard about these issues reject SI, presumably because they believe it to face insuperable objections. In

this case the defender of the EAAN may still be warranted in holding that $P(R/N\&E)$ is low even if they take it that $P(R/N\&E\&SI)$ is high. The argument would lose some of its scope, but, perhaps, none of its force. Naturalists, on the other hand, may find themselves more favourably disposed towards semantic inferentialism, were they to hold that it constituted a response (perhaps the only compelling one) to the EAAN. Even this simple response may not be straightforward. All that is required for the challenge of SI is the ‘weak’ variety in which having inferential relations of such and such a sort is *necessary* for having a content of such and such a sort. On weak semantic inferentialism (1**) remains unavailable as a justification for (1*).

3.2 Improbable?

It may be thought that SI *itself* is simply unlikely given naturalism and evolution. In this case the total probability of $(R/N\&E)$ would have to take into account the probability of R given $N\&E\&SI$ and the probability of R given $N\&E\&\neg SI$, which would be weighted by the probabilities of SI given $N\&E$ and $\neg SI$ given $N\&E$. Hence the total probability of reliability given naturalism and evolution would be equal to:

$$[P(R/N\&E\&SI) \times P(SI/N\&E)] + [P(R/N\&E\&\neg SI) \times P(\neg SI/N\&E)]$$

But since SI is an *a priori* thesis stating a (purported) conceptually necessary truth, $P(SI/N\&E)$ would be equal to 1 and $P(R/N\&E\&\neg SI) \times P(\neg SI/N\&E)$ would not be well defined. This strategy would make no advance; the total $P(R/N\&E)$ would have to be taken as equal to $P(R/N\&E\&SI)$.

3.3 Incompatible?

A more promising route would be to deny that SI *could* have a bearing on $P(R/N\&E)$ because SI is itself incompatible with materialism. On the materialist picture there are, presumably, either physical laws or metaphysical necessities governing what contents would be associated with which NP structures, and hence determining the function $f: NP \mapsto C$. In this case there *is* a lawlike connection between NP properties and belief contents, but no *prima facie* reason to think that this would map *appropriate* contents to roles. Given what we know about materialism, it may turn out to be the case that *every* content determined by an NP structure is ‘ $1+2=7$ ’. If this were so it would be a metaphysical necessity that every belief possessed by anyone was false, and hence the metaphysical probability for any given belief that it is true would be 0. One ought then to think of the EAAN as making a claim about *epistemic* probability; on naturalism and evolution, for any given belief, the epistemic probability that it is true is .5.⁶

But now SI looks to be incompatible with materialism, or, more carefully, it is epistemically probable that SI is incompatible with materialism. For SI holds that a belief has its content in virtue of the inferential role it plays, but it is epistemically probable that a final materialist theory of content will not associate contents with belief in this way. If one accepts SI *and* the conclusion that materialism entails that there need not be any principled connection between a belief’s role and its content then one is already rationally obliged to give up materialism. Plantinga however intends the EAAN to apply equally to naturalistic dualism. On naturalistic dualism there is a function $h: M \mapsto C$ from states of non-material minds to contents, and again the same considerations apply. Plantinga holds that it is unlikely that h will associate true beliefs with adaptive causal roles:

It isn't as if, as in theism, the person in question has been created in the image of a God one of whose outstanding characteristics is knowledge and understanding. [Plantinga 2008: 50]

But semantic inferentialism undermines (1**) in the dualist case for exactly the same reasons as in the materialist case. As such, even if SI is incompatible with materialism on naturalistic dualism the motivation to think that $P(R/N\&E)$ is low is still undercut.

3.4 Inconclusive?

Even though (1**) is *sufficient* for (1*) it may not be *necessary* for (1*). If the defender of the EAAN could motivate (1*) whilst accepting that (1**) is false, then the EAAN would be on very strong ground indeed. This task though would not be straightforward, for there is some reason to believe, on the SI picture of content, that our interactions with the world—indeed the kind of interactions that are required for behaviour to be *adaptive*—will structure the contents of our beliefs in such a way that they will have a good chance of being true:

The way the world is, constrains proprieties of inferential, doxastic, and practical commitment in a straightforward way from within those practices. So if I perceive a liquid as tasting sour, infer that it is an acid, infer further that it will therefore turn litmus paper red, and, intending to match a red pigment sample, accordingly dip litmus paper in the liquid, I may nonetheless subsequently acquire perceptually a commitment to the result being a blue, rather than red, piece of paper, and hence an acknowledgement of my practical failure. In this way I can find myself with incompatible commitments (which need to be sorted out if I am to remain entitled to any of my commitments in the vicinity). The possibility of incompatible commitments arising from the cycle of perception, inference, action, and perception reflects the way the normative structure of perception and action incorporates elements of the causal order. As a result, empirical and practical constraints get built into what commitments (including inferential commitments) one can sustain entitlement to. [Brandom 1994: 332-3]

To defend the EAAN in this manner, it must be argued that there are many different ways to behave adaptively, many different networks of belief whose content can be thus constrained by the cycle of perception, inference, action and perception, but which still differ in widespread and significant ways.

In a different context Angus Menuge makes a suggestion that may be adopted to motivate a claim of this sort. Menuge draws an analogy between theoretical instrumentalism and psychological instrumentalism. The leading idea is this: just as there are indefinitely many theories that save the phenomena—that account for the data—so too there are indefinitely many networks of belief that ‘save the phenomena’—i.e. that produce the adaptive behaviours in response to sense perceptions. Moreover, just as most instrumentally or empirically adequate theories are largely false, so too most psychologically or behaviourally adequate networks of belief are largely false.

[G]iven that there are a finite number of data points (stimuli and responses), and supposing that these exhaust the stimuli that a creature can observe and the responses it can produce, the basic lesson of Theoretical Instrumentalism remains. There are vastly more false mental model

that cover these data points than there are accurate (or nearly accurate) ones. Therefore it is much more likely that psychologically instrumentalist models will be available for selection than psychologically realist ones. [Menuge 2003: 161]

Matters are not straightforward however. Note that if SI is true, it is not the case that there are different belief systems compatible with *exactly* the same inputs and behavioural responses. This is because for any belief, no matter how recondite, there are some scenarios in which it can influence behaviour. Those who believe that paradoxes of self-reference can be captured by the Inclosure Schema behave differently, in the seminar room, than those who do not. Given SI, there is no way to sunder beliefs into those that pertain to action and perception and those that do not, as one might divide the entities quantified over by a theory into those that are observable and those that are not. As such, it is not the case that two distinct networks of belief could be behaviourally *identical*.

Nevertheless, a range of different networks of belief *could* be sufficiently behaviourally felicitous to have survived Darwinian pressures. The question then becomes whether, given SI+N+E, we are likely to have beliefs that, whilst behaviourally felicitous, are cognitively disastrous. It is certainly the case that a network of beliefs could be predominantly false whilst possessing these properties. Again in a different context, Plantinga considers a *naturalistic* scenario in which cognitive agents who describe everything using Russellian definite descriptions of the form ‘the creature, such that’ where creaturehood entails being created by God. On seeing a tree with red foliage, instead of believing ‘The tree has red foliage’ they believe ‘The creature such that it is a tree has red foliage’. These agents’ beliefs are largely false (remember, in this scenario, naturalism is true), yet adaptive and compatible with our inferentialist considerations. Yet, whatever the proportion of the beliefs these agents possess are false, it would be wrong to think of this as a *cognitive disaster*; as a really damaging kind of scepticism. In fact, it is not importantly different from situations we may well take ourselves to be in. Perhaps, for example, mereological nihilism is true, and macro-objects do not really exist. In this case, a huge proportion of our beliefs will be (strictly) false, just as a huge proportion of the beliefs held by the denizens of Plantinga’s scenario are false. But we do not, and ought not, think of this as a *pernicious* sceptical scenario akin to deception by an evil demon or ensconcement of one’s brain in a vat. It is incumbent on the defender of the EAAN to show that, given naturalism, cognitive disaster is likely.

The difficulty here is, as Plantinga’s example illustrates, that cognitive disaster, and the avoidance of cognitive disaster, cannot be simply parsed in terms of the proportion of one’s beliefs which are true or false. A different way to think about avoiding cognitive disaster, in the spirit of SI, is to think of our beliefs as embodying an understanding of the world. The agents in Plantinga’s scenario possess an (imperfect) understanding of their world that is not captured in terms of the truth or falsity of their beliefs, but can plausibly be captured in terms of their grip of the modal/nomic/dispositional structure of aspects of their world. Our interactions with the world, through the cycle of perception, inference, action and perception, are able to structure the content of our beliefs. When this process goes aright, the normative structure of one’s beliefs—the network of commitments and entitlements, what one is entitled or committed to infer from what—tracks the modal/nomic/dispositional structure of the world. This is not to say that our minds become a *mirror* of nature, but rather that they embody an *understanding* of the world by tracking (some of) its modal/nomic/dispositional properties.

Why is this significant? According to SI, the content of one’s beliefs is articulated

by its inferential structure. Given SI then, grasping the modal/nomic/dispositional character of objects or collections of objects *just is* having a network of beliefs whose normative structure tracks the modal/nomic/dispositional of the world. Yet, if one's network of beliefs is constrained by cycles of perception, inference, action and perception then it is probable, and perhaps necessary, that it will possess a normative structure that tracks aspects of the modal/nomic structure of the world to at least some significant extent. If then avoiding cognitive disaster is a case of grasping the modal/nomic/dispositional character of objects or collections of objects, then cycles of perception, inference, action and perception will structure one's beliefs in such a way that they avoid cognitive disaster. To put things the other way around: if a network of beliefs does not possess an inferential structure that tracks the modal/nomic/dispositional structure of the world to at least some significant extent, then it will not be behaviourally felicitous. By contraposition then, if a network of beliefs *is* behaviourally felicitous then it will possess an inferential structure that does track the modal/nomic/dispositional of the world to at least some significant extent. On SI, this is what possessing an understanding of the world would consist in. In brief, tracking the counterfactual nature of things around us is necessary for behavioural felicitousness and sufficient to avoid cognitive disaster. There is then, at least *prima facie* reason to think that (1) would be difficult to motivate, given SI; but the preceding account of cognitive disaster and its relationship to behavioural felicitousness is merely a sketch.⁷ Settling the issue would involve wading further into the deep and difficult territory of what constitutes a cognitive disaster. The challenge semantic inferentialism poses is the demand that such a task be undertaken.

Notes

¹The argument was originally presented in Plantinga [1991], and has since undertaken a number of refinements; cf. esp. Plantinga [1993, 2000, 2002, 2008, 2012].

²In what follows I will talk about NP properties 'determining' contents to denote both the reductive claim that NP properties are contents and the non-reductive claim that contents supervene on NP properties.

³Cf. esp. Dretske [1988, 1995].

⁴Cf. esp. Millikan [1984, 1989].

⁵In this way SI differs from other broadly functionalist theories of content. 'Generic functionalism', as Plantinga calls it, is the claim that what makes a neural structure or state *N* a belief *and* a belief with a particular content is that *N* plays a particular causal role. On generic functionalism, instead of there being a function *f*: NP $\mapsto$ C from Boolean combinations of sufficiently complex NP properties to contents, there is a function *g*: R $\mapsto$ C from causal roles to contents. As Plantinga [2011: 452] points out, this 'doesn't seem to place any constraints at all on the content enjoyed by [the neural structure that exemplifies that causal role]'. There is nothing to say that *g* associates adaptive roles with true contents; that further claim is not part of functionalism *per se*. Plantinga's response however cannot be brought to bear on SI, since according to SI the role a belief plays *constitutes* its content.

⁶Plantinga in fact takes the EAAN to be making a claim about metaphysical probabilities; yet, because of the way Plantinga understands materialism, we cannot coherently understand the argument this way. One option here is adopt a materialism that involves *contingent* psychophysical laws rather than reduction or strong supervenience. Plantinga however notes that the EAAN can be carried out in terms of epistemic possibility. What we are considering here is that for it to apply to all forms of materialism, the EAAN not only can but *must* be carried out in terms of epistemic possibility.

⁷Some authors (Crisp [2011], Otte [2011] and Rea [2002]) have suggested a weaker version of the EAAN which applies, not to all beliefs, but to some proper subset of them. The sketch above pertained most obviously to quotidian beliefs about what Austin called "medium-sized dry goods", and it is not immediately obvious whether or not it could extend to *recherché* metaphysical beliefs. It may be that whilst P(R/N&E) is not low *in general* given SI, it is low for beliefs that are not so immediately subject to the constraints imposed by cycles of perception, inference, action and perception; e.g. *recherché* metaphysical beliefs. If naturalism is to be counted amongst one's

recherché metaphysical beliefs, then the EAAN would still succeed. To motivate this, the defender of the EAAN would have to contend that our inferential practices are not well suited to matters metaphysical. Note that given Plantinga's motivation for (1), there is no reason to draw a distinction between *recherché* metaphysical beliefs and quotidian ones, but that, given SI, it may be the case that a restricted version of (1) could be motivated whilst (1) could not.

Works Cited

- Brandom, R. (1994), *Making It Explicit*, Harvard University Press.
- (2000), *Articulating Reasons: An Introduction to Inferentialism*, Harvard University Press.
- (2008), *Between Saying and Doing: Towards an Analytic Pragmatism*, OUP.
- Clark, K., and Van Arragon, R. (eds.) (2011), *Evidence and Religious Belief*, New York: OUP.
- Crisp, T. (2011), 'An Evolutionary Objection to the Argument from Evil', in Clark, K., and Van Arragon, R. (eds.) *Evidence and Religious Belief*, New York: OUP.
- Dretske, F. (1988), *Explaining Behavior*, Cambridge, MA: Bradford, MIT.
- (1995), *Naturalizing the Mind*, Cambridge, MA: MIT Press.
- Menuge, A., (2003), 'Beyond Skinnerian Creatures: A Defense of the Lewis/Plantinga Critique of Evolutionary Naturalism,' *Philosophia Christi*, 5(1): 143-165.
- Millikan, R., (1984), *Language, Thought and Other Biological Categories*, Cambridge, MA: MIT Press.
- (1989), 'Biosemantics', in *Journal of Philosophy*, 86: 281-97.
- Otte, R. (2011), 'Conditional Probabilities in Plantinga's Argument' in Clark, K., and Van Arragon, R. (eds.) *Evidence and Religious Belief*, New York: OUP.
- Plantinga, A. (1991), 'An Evolutionary Argument Against Naturalism', *Logos* 12.
- (1993), *Warrant and Proper Function*, New York: OUP.
- (2000), *Warranted Christian Belief*, New York: OUP.
- (2002), 'Introduction', pp. 1ff and 'Reply to Beilby's Cohorts', pp. 204ff, in Beilby, J. (ed.), *Naturalism Defeated?*, Ithaca: Cornell University Press.
- (2008), with Tooley, M., *Knowledge of God*, New York: Blackwell Publishing.
- (2011), 'Content and Natural Selection', *Philosophy and Phenomenological Research*, 83:2, pp. 435-458.
- (2012), *Where the Conflict Really Lies*, OUP USA.
- Rea, M. (2002), *World Without Design* (Oxford: Clarendon Press)
- Wanderer, J. (2008), *Robert Brandom*, Acumen.