

## A Tribute to Karen Neander (1954-2020)

Christopher Hill (Brown University) & Carlotta Pavese (Cornell University)<sup>1</sup>

It is 1980 when Karen Neander delivers a talk titled “Teleology in Biology.” She is a 26-year-old La Trobe graduate student speaking at the New Zealand division of the Australasian Association of Philosophy. The article, that would later become chapter three of her PhD dissertation (1984), advances and defends a novel, and soon extremely influential, theory of biological functions — the “selected effects” view — and with it lays out the foundations for Karen’s life long research project: marrying psychology with biology for the purpose of the naturalization of the mind.

The first part of “Teleology in Biology” argues for the centrality of the concept of function in biology, both in concept formation and in functional analysis. The second part develops a novel etiological account of functions. At the time of Karen’s writing, the most prominent theory of biological functions was the *goal account* (e.g., Ruse 1971, Wimsatt 1972, Baublys 1975, Boorse 1977). On this account, functions are contributions to goals, which, in physiology, are those biologically driven of survival and reproduction. According to the goal account, the function of the heart, for example, is to pump blood because that is how it typically contributes to survival and production. Karen argues that the goal account cannot properly deliver some crucial distinctions in biology — such as the distinction between *proper functioning* and *fortuitous effects* of biological items that are species-typical. To illustrate, the bridge of the nose is adaptive, for it helps keep up our spectacles. But its function is clearly not that of keeping up our spectacles. Rather, this is a fortuitous and species-typical effect that the bridge of the nose has.

An alternative to the goal account of functions was countenanced by Wright (1976, *Teleological Explanations*). Wright’s theory of functions was ‘etiological’ because it identified a function of something with its effects. In particular, according to Wright, the function of  $x$  is  $z$  if and only if a)  $z$  is the consequence of  $x$  being there and b)  $x$  is there because it does  $z$ . But Wright’s theory faced several problems. Perhaps the most glaring one is that, as Boorse (1976) had pointed out, this definition overgenerates, for it predicts that, for example, ‘obesity’ would have the function of promoting a sedentary lifestyle because that explains its continued existence. Already in this early dissertation chapter (p. 103), Karen shows how to revise Wright’s account in order to overcome this and other objections. Karen observes that if we explicitly restrict

---

<sup>1</sup> Thanks to Josh Armstrong, Alessandro Gatti, Justin Garson, Mark Gersovitz, and Alejandro Vesga for helpful suggestions and discussion, as well as to everybody who contributed to this tribute with quotes and remembrances of Karen.

functions to *evolutionarily selected effects*, then we can avoid the Boorse-style counterexamples. Obesity has not been selected by evolution because it promotes a sedentary lifestyle. The resulting account is simple and elegant: the function of an item  $i$  in an environment  $e$  is to do C if and only if  $i$  was selected in  $e$  because it does C.

The differences between this etiological account and the goal account of functions are subtle but critical. According to the goal account, a function is *adaptive*: it contributes to the goal of evolution, survival and reproduction. According to Karen's account, instead, a function is an *adaptation*: a function is an effect for which a trait was selected. While an effect can be adaptive, it might not be an adaptation, for it might contribute to survival and reproduction without having been developed, selected, and promoted by natural selection. And while a trait can be an adaptation, it does not need to be *currently* adaptive, as cases of maladaptations illustrate. Because of these crucial differences, the selected effects view of functions fares better than the goal account when it comes to delivering important biological distinctions such as that between proper functioning and fortuitous effects. It correctly predicts that the function of the bridge of the nose is not to keep up our spectacles, for that is not what it was selected for.

This work defending and developing a novel etiological conception of functions was presented at the AAP four years before Karen defended her PhD dissertation ("Abnormal Psychobiology," advised by Kim Sterelny, submitted in 1983 but defended in 1984), which was written almost entirely while on teaching duties at University of Sydney (1980-1983) and University of Wollongong (1983-1984). And it was eleven years after her AAP's talk that her views on functions landed a journal publication ("Functions as Selected Effects: The Conceptual Analysis Defence" *Philosophy of Science* 58 (1991): 168-184, and "The Teleological Notion of 'Function'", *Australasian Journal of Philosophy* 68 (1991): 454-468). Although these ideas only reached publication eleven years after their first presentation, they had been influencing philosophers around the world through the circulation of her unpublished dissertation work as early as 1980. David Papineau (King's College London) recalls, "I first met Karen back in the 1970s. In 1984 I examined her PhD 'Abnormal Psychobiology'. This was a wonderful piece of work. Amazingly for something by an unknown Australian research student, it was widely circulated, admired and cited long before she published anything."

The selected effects theory of functions is a core commitment of Karen's thought and will figure prominently in all of her work — from her dissertation, to her conception of the relation between psychology and biology, and her more mature theory of content and intentionality. The research project known as teleosemantics, of which Karen will be one of the main proponents, is just one of its applications. In Karen's dissertation, the selected effects theory of functions

was put to the service of a discussion of the ‘anti-psychiatry debate’, originated in the 60s and well alive at the end of the 70s and in the early 80s, when Karen was a graduate student. The debate concerned the status of psychiatry as a branch of medicine and the status of mental illnesses as *bona fide* medical illnesses. The side attacking psychiatry would question its status as a medical science, on the ground that many mental illnesses are, in the jargon, ‘purely functional’: irreducibly psychological and — supposedly — without organic basis. Karen proposes to look at whether psychiatry is a branch of medicine by answering a seemingly different question: whether biology and psychology are *continuous sciences*, and what the prospects are for an ‘abnormal psychobiology’ — a study of mental illnesses as biological dysfunctions. Karen argues that, *if* we are to understand mental illnesses as medical illnesses, we should ask if they are dysfunctions, in the biological sense of function. If they are, then psychology can be shown continuous with biology, for a core biological concept would be central in characterizing psychological concepts. In turn, Karen points out that psychology could be shown continuous with biology, *even though* psychology is ultimately irreducible to biology. Irreducibility does not entail discontinuity: for example, physiology, concerning the functional organization of the body, is irreducible to but continuous with molecular biology.

In this way, Karen turns the question: *is psychiatry a branch of medical science?* into the question: *can there be psychological functions and dysfunctions?* Karen responds that it is plausible to think that there *are* psychological functions and dysfunctions, for brains have evolved as they have partly because of the psychological capacities with which they endow us. She goes further: she argues that any plausible theory of mind is compatible with there being psychological functions in the strictly biological sense, and that *biological functionalism* — the view according to which mental states and processes can be characterized functionally and analyzed by way of functional analysis in the same way as physiological processes can be characterized functionally and analyzed by way of functional analysis (p. 177) — is the most promising form of functionalism. This part of the dissertation (pp. 180-192) contains, *inter alia*, a superb defense of biological functionalism against Block’s attacks on it in his “Trouble with functionalism” (1978).

Karen published little of this work on the anti-psychiatry debate (the interested reader can learn more about her views on mental illnesses in her “Mental illness, concept of” *Routledge Encyclopedia of Philosophy* (1998)). Yet, her defense of the viability of a psychobiology remains an incredible accomplishment. With it, Karen would set a model of how to do empirically informed work in philosophy to which she will live up throughout her career. While in constant dialogue with the sciences, she would never lose sight of the

distinctive contributions that philosophy can make to the deepest foundational questions.

Impressively, Karen's first publication in the philosophy of science is not in philosophy of psychiatry, and even predates her two papers on the concept of biological function published in 1991. Rather, her first published paper in philosophy of biology is a formidable attack against a prominent view on the explanatory role of natural selection. In 1988, while a lecturer at the University of Wollongong, her article "What Does Natural Selection Explain? Correction to Sober" appears in *Philosophy of Science*. The target is a view defended *inter alia* by Elliott Sober, at the time already one of most prominent philosophers of biology in the USA, who in his just published book *The Nature of Selection* (1984), had defended a 'negative view' of the explanatory role of natural selection. About this exchange with Sober, Justin Garson (Hunter College and The Graduate Center, CUNY) remarks: "One of the admirable things about Karen is that, despite her gentle demeanor, she did not back down from a fight. She even sparred with Elliott Sober on what natural selection explains. He is not someone I'd be eager to pick a fight with on what natural selection explains. But she did. What is more, at least for several commentators, she was right. Today, her view about natural selection is called 'the positive view' (in contrast to 'the negative view'), and it is endorsed by Peter Godfrey-Smith, Ruth Millikan, Patrick Forber, Jonathan Birch, and others."

While we cannot do full justice to the exchange between Sober and Karen in this short tribute, it is nonetheless worth highlighting some of Karen's key moves in it, for they constitute philosophy of science at its best. According to the 'negative view' of what natural selection explains, embraced by several prominent philosophers of biology such as Cummins, Sober, and Dretske, natural selection by itself can only explain the distribution of a trait in a population — e.g., why koalas have pouches. But it only plays a negative role in explaining why a *particular* koala has a pouch. To give a vivid illustration of the view, consider the following simple example of natural selection. Suppose in a generation, an organism produces two offspring and a mutation occurs, so that one of the individuals fails to resemble the parent in some respect: while one inherits trait B, the other does not. Suppose the new trait, G, and the old one, B, differ in promoting fitness, and as a result the old trait is selected against and disappears from the population. So, members of the following generation III — call them *a* and *b* — all have Gs. In this scenario, it seems natural to say that, while natural selection explains why all members of III have trait G, it does not explain why *a* and *b* each have G. What explains why *a* and *b* each have G is that *a* and *b* have inherited G from their parents, who also had G. And what explains in turn that their parents had G is that a mutation has happened. The mutation, along with the laws of heredity, does most of the explanatory work, while the selection itself only

helps get rid of any leftover individuals with the old trait B. In a slogan — which Karen herself made famous in her article “Pruning the tree of life” (1995) — according to this negative view, selection by itself *only prunes the tree of life*; it does not create new forms of life.

In a series of papers (“What Does Natural Selection Explain? Correction to Sober” *Philosophy of Science* 55 (1988): 422-426; “Pruning the Tree of Life,” *British Journal for the Philosophy of Science*,” 46 (1995): 59-80; “Explaining Complex Adaptations: A Reply to Sober’s ‘Reply to Neander’” *British Journal for the Philosophy of Science* 46 (1995): 583-587), Karen criticizes this negative view of the explanatory role of selection and develops a more positive view. In particular, contra Sober, she argues that natural selection can increase the probability that a specific gene sequence will arise in a population, and in that sense, natural selection can play into a positive explanation of why an individual has a certain trait. We can see this by imagining a hypothetical gene combination  $\langle A_2, B_2 \rangle$ , that codes for a useful trait. Suppose every member in the population has  $\langle A_1, B_1 \rangle$ . Now suppose that, by random mutation, a member gets  $\langle A_2, B_1 \rangle$ . If  $\langle A_2, B_1 \rangle$  is fitter than  $\langle A_1, B_1 \rangle$ , then natural selection will (*ceteris paribus*) tend to make it more common. But by making it more common it will increase the odds that a member of the population will come to have  $\langle A_2, B_2 \rangle$ . This example illustrates that, although there is the random generation of variation by means of mutation, it does not follow that only mutation explains the creation of genetic sequences (and the adaptations they give rise to). It might look that way if one only looks at *one* isolated mutation/select sequence. But millions of such sequences are involved in producing adaptations like our opposable thumbs, our eyes and ears, and the rooster’s crow. And these sequences are not causally isolated from each other. If that is right, selection does not just operate as a ‘sieve’: it can play a creative role as well.

The last exchange between Sober and Karen dates 1995. Meanwhile, a lot has happened in Karen’s career. After one year as a lecturer at the University of Wollongong (1987-1988), she had left this permanent position to accept a 7-years long post-doc position at ANU. She will be at ANU until 1995, when she will be leaving for a career in the USA. Karen always remembered this intense and fruitful period of research in her native Australia as one of the most rewarding times of her philosophical career.

1995 is also the year Karen’s first article on teleosemantics is published — “Misrepresenting and Malfunctioning,” (*Philosophical Studies* 79 (1995): 109-141) — which soon became a classic in the literature about naturalizing mental content. Together with “Dretske’s Innate Modesty” (*Australasian Journal of Philosophy* 74 (1996): 258-274), “Misrepresenting and Malfunctioning” inaugurates Karen’s application of the selected effects theory of functions to the study of intentionality, which will be a dominant theme in Karen’s thinking

throughout the following twenty-five years. A long tradition in the philosophy of representation takes it to be constitutive of representation that it can *mis*represent. This normative aspect of representation had been long thought to constitute a hurdle to the naturalization of this notion, as it seemed dubious that it can be reduced to naturalistically acceptable notions. Teleosemantics proposes to understand the normative aspect of representation in terms of the normative notion of a function and the distinction between proper functioning and malfunctioning. (For a careful and instructive, but not widely known, discussion of the relation between the normativity of functions and that of meaning, the reader might be interested in Karen's unpublished manuscript "The narrow and the normative" available on her Philpapers account).

A prominent objection to this general approach to the problem of intentionality is that, due to the intrinsic indeterminacy of biological functions, teleosemantic projects are affected by a corresponding problem of indeterminacy. A version of this objection is traceable to Fodor, who asked us to consider a frog that snaps at anything that is suitably small, dark and moving and thereby feeds itself. According to Fodor, if it was adaptive for the frog to snap at flies then it was equally adaptive for it to snap at small, dark, moving things, on the simplifying assumption that flies and small, dark, moving things were reliably co-extensive in the frog's natural habitat. We can equally well say that the function of the frog's detection device is to detect flies *and* to detect small, dark, moving things. So, if we try to determine the content of the representation by reference to the function of the detection mechanism, the content remains indeterminate between these two functions. Some had appealed to selected functions to overcome similar indeterminacy problems (Sterelny 1990, Millikan 1991). But as Karen points out, the appeal to selected functions by itself does not suffice to disambiguate content in this case. For in the case of the frog's detection device, its detecting small, dark, moving things and its helping the frog to catch and swallow something nutritious *both* played a causal role in selection of the relevant representation producing or consuming systems: for it was by detecting small, dark moving things that the frog got fed.

In "Misrepresentation and malfunctioning," Karen defends a novel approach to teleosemantics (a 'low church' approach, as she labels it) which affords a novel and principled solution to this indeterminacy problem. Karen first argues (p. 116) that plausible attributions of functions to an organism and its parts have a place in a functional analysis which decomposes the organism into simpler and simpler parts and describes the contributions of each part to the overall activities of the organism. If so, then when asking about the function of a trait, we might be asking about the function(s) of the whole organism to which the trait contributes; or we might be asking about the function that is most specific to that trait. So, for example, the function most specific to the heart is to pump blood. But

at the highest level of functional analysis, the description of the heart's functions might include the functions of the whole organism to which the heart contributes, including reproduction and genes proliferation.

Now, Karen points out that a 'high church' teleosemantics gives priority to a *higher level* of analysis. So, for example, according to a high church approach, in the case of the frog, the function relevant for determining the content of the frog's representation might be *to help catch something nutritious*. By contrast, according to a 'low church' teleosemantics, in choosing among the relevant content, we should give priority to the function that is most specific to that trait, rather than to the functions of the organism to which that trait simply contributes. In Fodor's example of the frog, the lowest level of functional analysis is *detecting small, dark, moving things*. For that is the function most specific to the frog's detection device and for it is *by* detecting small, dark, moving things that the frog detects frog-food, feeds, and can reproduce and proliferate its genes (p.130). This distinctive approach to teleosemantics, characterized by a more mechanistic teleology, more sensitive to the demands of a detailed functional and cognitive analysis and so conducive to more psychologically plausible verdicts, will become central to Karen's view in *A Mark of the Mental* (2017) but is already clearly there in this earlier formulation of her theory.

From 1996 to 2002, Karen takes up positions first as an assistant professor and then as an associate professor in philosophy and cognitive science at Johns Hopkins University. To this period belongs a defense of teleosemantics against the Swampman's objection ("Swampman Meets Swampcow," *Mind and Language* 11 (1996): 118-129), which we will discuss in more detail below; a critical discussion of Godfrey-Smith's environmental complexity thesis ("The Function of Cognition: Godfrey-Smith's Environmental Complexity Thesis", *Biology & Philosophy* 12 (1997): 567-580); and an article raising a novel objection to representationalism about consciousness ("The Division of Phenomenal Labor: A Problem for Representational Theories of Consciousness", *Nous* 12 (1998): 411-434).

From 2002 to 2006, Karen is a professor of philosophy at UC Davis. And in 2006, she arrives at Duke, where she will remain until her premature death. During these years at Duke, she is a role model for many junior professors, who cherished her for advice, encouragement, as well as her refreshing sense of humour — not to speak of her brilliance, the incisiveness of her comments, and the depth of her philosophical commitments. In her remembrance, Sara Bernstein, a colleague of Karen's for several years at Duke and now at Notre Dame, writes: "She was a brilliant, top-notch philosopher: thoughtful, careful, and incisive, with no trace of ego in her work. She was interested in developing good ideas over the long haul. She did great philosophy, slowly, and with great care for the details. She didn't want to make things easier just because one could. She didn't like flash. She didn't go for philosophical trends just because they were trendy. She was unimpressed

with simple solutions. She didn't overstate her intellectual accomplishments, though they were formidable. She was in it for the ideas."

While at Duke, Karen co-authors with Alex Rosenberg two articles on functions, both defending and expanding on the etiological theory of functions. The first ("Are Homologies (Selected Effect or Causal Role) Function-Free?" in *Philosophy of Science*, 76.3 (2009): 307-334) defends the role of etiological functions in understanding "homology judgements" in biology, such as the judgements that the bones of the vertebrate forelimb are homologous — i.e., that their "similarity" is due to common descent. The second article ("Solving the Circularity Problem for Functions A Response to Nanay," *Journal of Philosophy*, 109.10 (2012): 613-622) defends the etiological theory of functions against an alleged circularity challenge and launches a devastating attack against a modal account of functions.

Most of her other work in this period is preparatory to her book. Noteworthy in this respect are "Content for Cognitive Science", in *Teleosemantics*, (Papineau and McDonald (eds.) OUP 2006); "Teleological Theories of Mental Content: Can Darwin Solve the Problem of Intentionality?" in the *Handbook of the Philosophy of Biology* (Ruse (ed.) OUP, 2008); her *Stanford Encyclopedia of Philosophy* entry, "Teleological Theories of Content" (originally published in 2004 and last revised in 2012); and "Toward an Informational Teleosemantics" in *Millikan and Her Critics* (Kingsbury and Ryder (eds.) Wiley Blackwell 2013). In her paper "Functionalist Analysis and Species Design" (*Synthese*, 194 (2017): 1147-1168), Karen holds that a concept of normal proper function is used in physiology and in cognitive science to explain how bodies and brains operate and argues that this notion is best understood in terms of etiological functions rather than in terms of Cummins's (1975) functions. This material will become the cornerstone of chapter 3 of *A Mark of the Mental*.

From this period, we would like to single out Karen's unpublished contribution to the 2015 Minds Online Conference, on *The Brains Blog*, titled "Why I'm not a Content Pragmatist." Here, Karen motivates with exceptional clarity her realism about mental representation and mental content. This article is a discussion of *content pragmatism* — a view defended by Frances Egan in "A Deflationary Account of Mental Representation", for which Karen was a respondent at an SSP meeting that took place at Duke University in June 2015. According to *content pragmatism*, content and mental representation are not real but they can nonetheless play a helpful role in psychological explanations. Although *content pragmatism* purports to be an alternative to *eliminativism*, according to which mental representation is not real and has no explanatory role to play, as well as an alternative to *dualism*, according to which mental representation is real but in principle inexplicable in naturalistically acceptable terms, Karen argues that *content pragmatism* is no genuine alternative to either



view. For according to *content pragmatism*, determinate content ascriptions are appropriate or inappropriate depending on what the *explanatory aims* are. However, explanatory aims are themselves intentional phenomena that *content pragmatism* leaves unexplained. So in this respect, *pragmatism* is very much like *dualism*. On the other hand, if *content pragmatism* amounts to the view that mental representation is currently explanatorily helpful but not in the long run and all things considered, then *content pragmatism* is not really an alternative to *eliminativism*: it is a version of *eliminativism*, but “albeit one that recommends patience (p. 4)”.

We have reached 2017 — the year in which *A Mark of the Mental* (MIT Press) is published. Peter Schulte, at University of Zurich, writes: “Karen Neander’s earlier papers on teleosemantics (e.g., “Misrepresenting & Malfunctioning”, 1995) already had a deep and lasting impact on the debate, but *A Mark of the Mental* firmly establishes her as one of the leading figures in the field, on equal footing with Ruth Millikan, Fred Dretske and David Papineau.” It is indeed hard to overstate the importance of Karen’s contribution in her *A Mark of the Mental*. It is to a discussion of this book that we turn next.

Karen’s brilliant book is concerned with the ancient and very difficult problem of intentionality — the ability of the mind to reach out into the world and grasp independently existing objects, properties, and states of affairs. One can see bookshelves across the room, hear a coyote howling outside, remember yesterday’s pizza, believe that Cleopatra was of Macedonian descent, and search for the largest prime number after 1,000. In all of these cases, one is in a state that is concerned with an item existing outside one’s mind. That is, one is in a mental state that refers to an external item. What does this mental reference, or aboutness, consist in? Karen’s book is an attempt to answer a significant part of this question — the part that is concerned with the intentionality of perception. In other words, it seeks to explain what it is for a perceptual state to be a perception *of* a specific external object or property or event. The attempt is innovative, rigorous, scientifically well grounded, brave in its willingness to recognize and address problems, and, we believe, largely successful. As a result, reading the book is an exciting intellectual adventure. It also provides the pleasures of discovery and enlightenment. At the same time, though, one reads it with a sense of loss and regret, for as one turns the pages, one keeps thinking that Karen would very likely have been able to extend our understanding of other parts of intentionality had she lived.

The book can be summarized briefly by saying that it is concerned with explaining intentionality in terms of the idea that mental states are representations — they are about items in the world because they represent items in the world. It is also concerned to explain representation in naturalistic terms — more specifically, in terms of the notion of a biological function and the notion of information.

Turning now to the more concrete themes of the book, it can be described as making five main proposals. The first of these is an etiological theory of functions that, as we have seen, Karen continued to develop and promote throughout her career. In the version that appears in the book, the etiological theory states that it is the function of an item X (where X is a type of state or process or mechanism) to do or cause Y if X was selected to do or cause Y. In the case in which X is a biological item, this proposal becomes the claim that it is the function of X to do or cause Y if X was selected by evolutionary processes because X did or caused Y on a number of historical occasions. This is the bare bones of the etiological theory, but Karen expands it to allow for the ways in which functions determined by natural selection can be modified or supplemented by learning processes in the histories of individual members of a species. Like evolutionary processes, learning *selects* behavioral and cognitive tendencies, the difference being just that learning operates at the level of the individual.

The second main proposal is that low level representations in biological systems, such as the representations that figure in perception, can be explained in terms of functions. More specifically, the idea is that a perceptual state X can be said to represent an entity Y if it is the function of X to encode information about Y. This is the defining principle of *teleosemantics*, a general approach to explaining representational content that was pioneered by Ruth Millikan, Fred Dretske, David Papineau, and Karen. Teleosemantics is attractive for a number of reasons, one of the most important of which is that it promises to explain the normativity of representation. Representations can *misrepresent* the items they refer to, and they can also *misrepresent* the world as containing items that it lacks, such as unicorns and phlogiston. That is, there is such a thing as representational error. But to describe a representation as erroneous is to evaluate it normatively, classifying it as a kind of failure. Teleosemantics maintains that it is possible to reduce this sort of normativity to malfunctioning — the failure of states or mechanisms to perform the functions that natural selection has conferred on them.

The third proposal concerns the nature of the functions that figure in perceptual representation. We have already seen that the relevant functions are informational functions, but what is information? Karen answers that information is grounded in causation: roughly, an item X carries information about another item Y to the extent that features of X were caused by features of Y. Thus, for example, tree rings carry information about the growth of trees because they are caused by the growth of trees. In combination with the central claim of teleosemantics, this view of information leads to the idea that a perceptual state R in a biological system S represents the property C just in case it is the etiological function of S to produce R when S is causally activated by an object with C. This is Karen's *simple starter theory of perceptual representation*. Note its clarity and simplicity!

The fourth proposal adds a new dimension to the simple starter theory. This is the idea that some systems of perceptual representation are analog in a sense that Karen explains. Suppose that the members of a system of representation  $R_1, R_2, \dots, R_n$  stand for a family of properties  $P_1, P_2, \dots, P_n$ . The system is analog in nature if there are similarities between  $R_1, R_2, \dots, R_n$  that correspond to similarities between  $P_1, P_2, \dots, P_n$ . To illustrate, levels of mercury in a thermometer constitute an analog representational system because differences in levels of mercury correspond to differences in the magnitudes of the temperatures that the levels of mercury represent. One can determine how close two represented temperatures are by checking to see how close the two levels of mercury that represent them are. In Karen's view, perceptual representations tend to be located in multi-dimensional similarity spaces that correspond to the multi-dimensional similarity spaces in which properties of objects are arranged, and distances between representations in any one dimension are correlated with distances among the represented properties in the corresponding dimension. Perceptual, cognitive, and motor systems exploit these structural correspondences between systems of representation and systems of properties in conducting their business.

This fourth proposal adds an altogether new dimension to teleosemantics, and so does the fifth. All theories of representation must answer the question of why perceptual representations represent distal objects like predators, prey, and fruit rather than more proximal entities like packets of light carrying information about distal entities, or retinal projections of such entities, or thalamic projections of the entities, etc. Most theories have been unable to meet this challenge but Karen's fifth proposal advances the discussion significantly. On her view, a perceptual representation  $R$  represents an entity  $X$  rather than an entity  $Y$ , even though both  $X$  and  $Y$  are causes of  $R$ , if it is true that  $X$  causes  $R$  by virtue of causing  $Y$ , and it is not true that  $Y$  causes  $R$  by virtue of causing  $X$ . In other words,  $R$  represents  $X$  rather than  $Y$  if  $X$  is the originating cause of the structured signal that eventually causes  $R$ , while  $Y$  is only a mediating cause. This is the intuitively correct result, because distal causes like predators and prey are precisely the causes of reflected light that originally impose the structural features on the light to which the relevant perceptual mechanisms are sensitive.

There is much more of great interest in the book than these five main proposals. For example, there is an extended argument, embedded in a richly informative chapter about the visual systems of toads, that the properties represented by perceptual states are low level properties like sizes, shapes, colors, locations, and directional velocities, as opposed to higher level properties like *face* and *angry face*. (If you like toads, this is the book for you.) There is also a careful and largely persuasive chapter maintaining that teleosemantics delivers unique contents for perceptual representations.

There is an important objection to teleosemantics that Karen does not address in the book, the so-called Swampman objection. Swampman is a human-like creature that comes into existence as a result of quantum fluctuations in swamp gas. He is endowed with a brain, internal organs, and sensory systems just like ours, and, as a result, he engages in the same sort of perceptual processing as we do, and he also engages in the same perceptually directed behaviors. As a result of these similarities, we are strongly disposed to think of Swampman as being in perceptual states that represent the external environment. But reflection shows that these representational states cannot be explained in teleosemantic terms. After all, Swampman is not a member of a species with an evolutionary history. So it seems that he is a counterexample to teleosemantics in general and to Karen's theory in particular.

We mention this problem because we want readers to be aware that Karen discusses it elsewhere and offers a compelling answer. (The locus is "Swampman Meets Swampcow," *Mind and Language* 11 (1996): 118-129.) The main point of her response is that it should be science, and not the philosophical imagination, that determines our beliefs about the essential natures of natural phenomena. Our ordinary concepts of natural phenomena often track what science reveals to be natural kinds, but these tracking relationships tend to be quite loose. Accordingly, while the philosophical imagination is a useful tool when one is concerned to chart the boundaries of the sets of actual and possible objects to which our ordinary concepts apply, it can be a poor guide to the boundaries of a natural kind that is imperfectly tracked by the concept. It is science, and not a priori reflection, that can chart the latter boundaries. Moreover, it is the latter boundaries that should matter when we are concerned to map the essential structure of reality. Metaphysics should be the story of what there is, not the story of how our everyday concepts represent what there is. Karen illustrates these points by asking us to imagine Swampcow, a creature that is exactly like a cow except that it comes into existence as a result of quantum fluctuations in swamp gas. Now contemporary evolutionary biology individuates species in terms of shared phylogenetic ancestry; that is, the boundaries of a species are determined by common descent. Because of this, Swampcow definitely does not count as a cow according to contemporary scientific standards. Could a philosopher sensibly maintain that these standards have been overturned because she has been able to imagine Swampcow? Karen playfully imagines a philosopher rushing into a meeting of evolutionary biologists to announce that she had found a counterexample to their theory of what makes a cow a cow. What sort of reception would she receive?

Karen's book has found a large and enthusiastic audience: it has been used as a text in seminars around the world; its ideas figure prominently in major new books by Nicholas Shea and J. Robert G. Williams; and it will soon be the topic of a

symposium in *Philosophy and Phenomenological Research*. It will no doubt continue to be a major guidepost in the literature on intentionality for many years to come.

We conclude this discussion of the book with a quotation from David Chalmers, one of the first philosophers to use it in a seminar: “Karen’s superb book has kicked off a burgeoning revival in the project of naturalizing intentionality since it was published in 2017. It shows a way forward with a simple and powerful version of teleosemantics for perceptual content. Now everyone is picking up on it by building their accounts of intentionality with Karen’s theory as a basis. It’s clear that her work will last.”

Towards the end of her life, Karen was working on the problem of explaining the sort of higher level intentionality that is characteristic of concepts and propositional attitudes. In particular, Karen took the next step of her project to be that of explaining the intentionality of ‘basic’ conceptual representations of which certain nonhuman animals are capable, before moving to studying the intentionality of more sophisticated forms of conceptual representation. On this topic, Karen held a discussion group at Duke in the Spring 2016. The last part of her last graduate seminar in the Fall 2017 was also devoted to a discussion of these issues. She was making progress, and she was able to give several talks that summarized this research. If only...

© 2020, Christopher Hill and Carlotta Pavese. All rights reserved