

Functional Classification and Dynamical Systems Theory

Michael Oliver

In "Meaning as Functional Classification," Wilfrid Sellars presents his perspective on the relation between syntax and semantics. In this piece, Sellars suggests that we ought to understand meaning as functional classification. In what follows, I will explicate Sellars's argument and its implications. I will then critically evaluate this Sellarsian account. In my analysis, we will find that Sellars has no doubt made some very pointed insights. However, I will argue that many of the key insights find a finer articulation in terms of dynamical systems theory (DST), rather than in terms of the symbol-processing tradition as Sellars may have preferred. In articulating Sellars's insights in terms of DST, we will see something like a symbiotic relationship. In other words, the synthesis of the Sellarsian tradition with DST yields a stronger account for both of the involved traditions, creating a complimentary whole.

Sellars begins his exploration of the relationship between syntax and semantics by articulating, and then endorsing, Gilbert Harman's three-tiered approach to meaning. This three-tiered approach is not entirely new, for it was conceived by reflecting on three different approaches to explaining what it is for a linguistic expression to have meaning. Originally, these approaches were believed to be in competition with one another as they were believed to be different approaches to the same question. However, as Harman argues, it appears as if these three approaches were not exactly in competition with one another because they seemed to be addressing different levels of the same question (Sellars 1974, p. 417). These three levels are:

1. Language at the level of speech/thought at the distinctly human level.
2. Language as the medium of communication between agents.
3. Language at the level of social practices such as acts of stating, promising, etc. (Sellars 1974, p. 417)

The argument, roughly formulated, is that these different approaches are not necessarily in competition with one another because they presuppose one another. For example, to talk about

meaning at the level of social practices, such as acts of stating, promising, etc., we need to first give an account of communication (Sellars 1974, p. 418). How can we articulate a promise if we are failing to communicate that promise to another agent? In short, we cannot.

After Sellars explicates Harman's three-tiered approach, he begins to construct his own account of language and meaning at the distinctively human level, that is, level one.¹ In laying out the tools for his discussion, Sellars leans on a notion that he calls "Verbal Behaviorism."² For Sellars, Verbal Behaviorism (VB) is the claim that what is described by a level one account is what it is to think at the characteristically human level (Sellars 1974, p. 418). After some preliminary qualifications, such as recognizing that his VB is not intended as a complete account of thinking, Sellars goes on to use VB in order to clarify some problems in language. In this sense, VB is a sort of specialized tool crafted for the purpose of explicating and clarifying issues that arise when giving a level one account of meaning. For Sellars's VB, there are two senses to having the thought occur to one 'that-p'. The first sense, which ought to be rather obvious for any tool that bears the name VB, is saying 'p'. The second sense, and probably more interesting for our purposes, is a short term propensity to say 'p' (Sellars 1974, p. 419). In addition, Sellars's VB is such that the thinking-out-loud 'that-fa' is the manifestation of some part of the conceptual function played by 'f' and some part played by 'a' within the PMese linguistic community.

The next tool that Sellars uses is termed "indirect discourse." Indirect discourse allows us to take expressions such as:

The thought that there is mold in the coffee mug occurred to Pat.

Then express it in more VB-friendly terms,

Pat says, 'There is mold in the coffee mug'.

or,

Pat had the short-term disposition to say 'There is mold in the coffee mug'.

This is possible because according to Sellars's VB, having the thought occur to one that 'p' cashes out the same if expressed as

¹ It is interesting to note that in giving his account, Sellars takes care to talk about how we can appreciate the virtues of Harman's three-tiered accounts even if we do not endorse something like "inner conceptual episodes." This will become increasingly more interesting when we attempt to articulate this account in terms of DST.

² This is not to be confused with Skinner's verbal behaviorism, though it is related in a number of complex ways.

something someone has said or had the disposition to say. This kind of ascription is where we can see a little bit of Sellars's commitment to "inner conceptual episodes" creeping into the picture. However, Sellars points out that we can be relatively agnostic about such conceptual episodes. From one perspective, probably most closely associated with the symbolic approach to cognition, we can view these ascriptions as picking out discrete mental states in the agent, where mental states can be cashed out in any number of ways. For example, mental states could be cashed out as meaning a functional state of the central nervous system, or, perhaps, as picking out some ontologically distinct category of things. For my purposes, I will be cashing these mental states out as being a theoretical enrichment of the largely behavioristic framework, but more on why we would do such a thing later.

Before we move on to a discussion explaining how we might hope to understand ascriptions of the sort mentioned above, we need to first qualify our use of VB. Sellars points out, after recognizing that these kinds of ascriptions can be useful, that these ascriptions are subject to a number of criticisms. One criticism is that this process can only be directed at a narrow selection of content. Another is that the process of indirect discourse fails to account for the fact that we could ascribe some measure of indirect discourse to some agent, and in some sense this could be done correctly, yet the actual utterance made by that agent could differ in a number of important ways. Both of these issues need to be addressed if we are to make any fruitful use of VB.

In an attempt to address these criticisms, Sellars develops a working thesis that will allow us to handle these problems. Sellars suggests that we understand the process of indirect discourse, at least in the context of VB, as giving a functional classification to something that someone says, or to describe what an utterance means (Sellars 1974, p. 421). This functional classification involves a specific formulation of expressions such that we can handle things as diverse as the syntax of an expression, the articulation of some sensory stimuli, or the connection between practical thinking and behavior. Understood as such, quoted expressions, such as the example above involving Pat, are ways of saying that an agent said, or had the short term propensity to say, something functionally equivalent to the ascribed content. For this expression to have meaning then, is for the agent to recognize the functional role fulfilled by the expression that they have uttered.³ The question

³ This insight is explored in greater depth in Sellars's paradigm-shifting essay "Empiricism and Philosophy of Mind," (1977) particularly sections 32-38.

immediately arises, “How does one come to recognize the functional roles fulfilled by an expression?” Before we continue reviewing Sellars’s thesis and how it might address the criticisms that have been presented, it seems necessary to address this question. In addressing this question, Sellars will explore pattern-governed behavior which will then lead me to introduce a few notions from the terminology of DST.

The short answer to the question as to how one might come to recognize the functional role of an expression is that it is the result of pattern-governed behavior (Sellars 1974, p. 423). However, this alone won’t suffice. What kind of pattern is it? Do we have any reason to think it really might be the result of pattern-governed behavior? There are at least two main kinds of patterned-governed behavior that we might see exhibited by language. The first of these patterns develops through the direct intention of some agent or central controller. This doesn’t seem to fit the profile of language as we experience it, since we see a variety of changes that occur in languages that do not appear to be the result of any particular intention. For example, some phonological changes involve the dropping of the latter phonemes in a word. This, in many cases, is not the result of a direct intention to drop such phonemes, but rather as a result of those phonemes gradually being unintentionally nixed since they weren’t playing any particularly important role in distinguishing words.⁴ The second kind of pattern, the kind that appears more compatible with meaning and language, is developed as a result of various propensities, propensities to reinforce behavior that conforms to the pattern and to reduce behavior that does not conform to the pattern. Acquiring the recognition of the functional classification of an expression is this latter kind of pattern-governed behavior

Within the latter category of pattern-governed behavior, we can make a further distinction. Borrowing some terminology from Andy Clark’s 1997 book, *Being There: Putting Brain, Body, and World Together Again*, we can make the further fruitful distinction between pattern-governed behavior that exhibits direct or indirect emergence. Direct emergence relies most heavily on the properties and organization of the individual components while giving only a slight nod in the direction of environmental conditions (Clark 1997). As example, Clark has us imagine a traffic jam with no environmental cause, such

⁴ We can also see languages resisting being driven by the intention of speakers as we might see in the history of languages such as Esperanto and the efforts of what Pinker refers to as the “Language Mavens” in *The Language Instinct*.

as an accident or debris on the road. In this traffic jam drivers are following two rules:

- I. If you see another car close ahead of you, slow down.
- II. If not 1, then speed up so long as you are under the speed limit.
(Clark 1997)

In the context of this example, the individual components are the drivers/vehicles, while the relations between the components is their location and distribution. While following these two simple rules, we see the emergence of the more complex phenomena, the traffic jam. The traffic jam emerges because the vehicles are speeding forward until they reach the next vehicle which causes them to slow down, which causes the vehicle behind them to slow down, etc. This example is illustrative of direct emergence because the phenomena emerges as a result of the two rules and relation between the individual components, not because of some other environmental circumstances. As an example of indirect emergence, we can take a look at the emergentist solution used by a group of researchers to have robotic insects collect pieces of wood into piles (Resnick 1994). This solution involved the use of two rules.

- I. If you are not carrying anything and you bump into another wood chip, pick it up.
- II. If you are carrying a wood chip and you bump into another wood chip, put down the wood chip you are carrying.(Clark 1997)

Again, we have two basic rules. The difference between this scenario and the traffic jam is that this case requires the interaction of the agent with the environment, where environment refers to that which is distinct from the individual components and their relations to one another. The environment in this case is the wood chips. The agents are the robotic insects and their rule set. In other words, the environment, is that which is not the rule-following, or pattern-governed, constitutive components. In this case, the researchers found that given these rules and some set of environmental conditions, e.g., the locations and distribution of the wood chips, the robots were able to stack the wood chips into a number of piles, fulfilling the goal of the experiment. Given this further distinction, is the pattern-governed behavior, as described by Sellars, more consistent with direct or indirect emergence? Is emergence exhibited at all?

If we are going to address this question adequately, we need to return to Sellars's essay. For Sellars, there are three kinds of pattern-governed linguistic behavior: "Intra-linguistic Moves," "Language

Entry Transitions,” and “Language Departure Transitions” (Sellars 1974, p. 424). Intra-linguistic moves are what we might think of as traversing the space of reasons, that is, making inferences both theoretical and practical. To begin with, Intra-linguistic Moves are clearly rule-governed since they involve some set of rules of inference. In addition, these rules are manifested as a result of patterned-governed behavior. It also appears to be the case that this pattern-governed behavior is of the second kind, meaning that it does not seem to come about because some agent intended it come about. We can find support for this claim by reflecting on the practice on logic. Logic appears to be the formalization of, or an attempt to make explicit, what people actually do when they reason.⁵ However, does this pattern-governed behavior give rise to any emergent properties? When we make inferences, we are participating in the pattern-governed behavior that is Intra-linguistic Moves. As such, our inferences are subject to evaluation. This means that when we make an inference we are in some sense leaning on other agents to either approve of, or disapprove of, our inferences. Since there doesn’t appear to be any important degree of non-agent environmental interaction, then this could be best classified as a pattern-governed behavior that results in direct emergence. Where the emergent properties arise, the behavior can be described as correct or incorrect, a feature for all of Sellars’s pattern-governed behaviors.

The next kind of pattern-governed behavior, necessary for any language, is Language Entry Transitions. These are inferences made from objects in the world using something like Sellarsian sense impression inference. This allows agents to respond to objects and potential inner states with the proper linguistic behavior (Sellars 1974, p. 423). Since interaction with the environment, again, some non-agent component, is required to lead to a linguistic response, this pattern-governed behavior can properly be described as resulting in an indirect emergence. In addition to be up for appraisal in terms of correctness and incorrectness, indirect emergence seems to have a few other interesting properties. If we are going to say that meaning is functional classification, then offering a functional classification of linguistic behavior that falls into this category necessarily involves some kind of response to objects in perceptual situations. Furthermore, since the classification involves, in some sense, a response to perceptual situations, this classification extends

⁵ I suppose it is conceivable that someone actually intended logic to be the way that it is and in doing logic what we are doing is tapping into this supernatural plan. However, I can’t seem to take this all that seriously since positing such a supernatural realm raises even more difficult questions that it answers.

out into the world, or, if you prefer, the environment. This means that if we are going to represent the meaning of such expressions, then our representation, in some sense, involves the world. The representations are not merely inner conceptual episodes, for they require an environment. This form of representation could not be properly manifested in a situation where an agent was detached from an environment. They are representations that straddle agent and world. In straddling agent and world, we need to resort to some other tool set to give them an adequate description. No longer can we focus merely on discrete, inner states. On the other hand, it could be the case that we don't represent the meanings at all, but rather exhibit the meaning. If this is the case, then it appears that a representation of a meaning is either not merely inner, or there is no real sense to talking about representations in the first place.

Finally, we have the last pattern-governed linguistic behavior. Sellars refers to this category as Language Departure Transitions. In these cases, the agent responds to a linguistic conceptual episode, such as, "I will now take a sip of from my mug," with the appropriate action, i.e., reaching for the mug, lifting it to one's mouth, etc. The interaction with the non-agent environment should be immediately apparent. Again, we have a behavior that could be described as correct or incorrect, but we also can point to a number of other interesting properties. If we are going to represent the meaning of an expression that belongs to this kind of pattern-governed behavior, then we have to lean on the environment once again. The functional classification could not be recognized unless one also represented the function of that expression in one's environment, i.e., the action following the conceptual episode.

The implications of these findings seem quite interesting. These findings suggest that if meaning really is some matter of functional classification, then we need to reevaluate our notion of a representation. It appears that there are roughly two paths we could take in coming to a new understanding. Our notion of a representation either needs to be capable of incorporating a representation straddling agent and world, or we need to dispose of the notion of a representation altogether, replacing it with a more sound notion. However, before we even attempt such a drastic course of action, we need to see if Sellars's VB, indirect discourse, and his notion of a pattern-governed behavior can even address the criticism laid out in the earlier sections of this piece.

As you may or may not recall, there were a couple of immediate objections raised to Sellars's notion of indirect discourse. The first was the idea that Sellars's indirect discourse seems rather narrow in scope. However, this worry is can be quickly suppressed by stating

that we are only attempting to give a level one account of meaning, that is, an account of meaning and thought at the distinctly human level. There is a great deal more behavior that can be properly described as intelligent, but this is not in the scope of the level one account. What is true though, is that indirect discourse seems to only be effective, at least as formulated up until this point, in the language of the speaker, in our case English. This objection is related to the second objection set forth above which stated that we could properly ascribe some VB to an agent, yet the agent uttered something non-trivially different from the ascribed VB. The solution to this problem lies in the formulation of our expressions, which will need some obvious refinements. Luckily, Sellars does most of this work for us. If you will recall, our primary formulation of indirect discourse landed us with an expression such as:

Pat says, 'fa'.

or, as Sellars points out, to abstract from the modality with

Pat 'fa's. (Sellars 1974, p. 426)

In this case we would be expressing the idea that Pat did 'f' and then did 'a'. In addition, we could understand 'fa' as being any number of sortal predicates which hold between between 'f' and 'a' (Sellars 1974, p. 427). But how does this resolve the issue of an agent saying something non-trivially different from the ascribed content? The solution is quite simple once we recognize that we are only using VB as a model to avoid any ontological concerns. In order to resolve this issue we simply need to make use of the famous (infamous?) Sellarsian block quotes. We can use the block quotes as a way of classifying some expression into some functional role. This works as long as we put block quotes around a classification with which a speaker is already familiar. Since I assume that the reader is a proficient English speaker, we can go ahead and specify our functional classifications in English. This would land us an expression like

Pat says ·fa·

Which would say something to the effect or Pat said or had the short term disposition to say something functionally equivalent to 'fa'. We can even make use of this format, after a little more refinement, for handling things like translations and abstract entities. For example, we can express the meaning of the German word 'und' by articulating its functional role with

'Und's (in German) are ·and·s. (Sellars 1974, p. 431)

This expression articulates the idea that ‘und’ (in German) plays the same functional role as ‘and’ in English, namely conjunction. Furthermore, we have a nice tidy way of handling abstract entities, without having to commit to a supernatural ontology, with expressions like

‘Dreieckig’ (in German) stands for triangularity. (Sellars 1974, p. 432)

which at first glance seems to be saying that triangularity is a name referring to a non-linguistic entity, a curiosity since it appears to express a relation between a linguistic and non-linguistic entity. However, once we put it into the format of functional classification that we have been working with, we see that these confusions quickly dissipate.

‘Dreieckig’ (in German) stands for triangularity.

becomes

‘Dreieckig’s (in German) are ·triangular·s. (Sellars 1974, p. 433)

Presented as such we begin to see more clearly the relationship between the two terms, this relationship is that they place the same functional role. We are functionally classifying ‘Dreieckig’s into the category of ·triangular·s.

With these refinements in place we see the objection that an ascription of VB to some agent could be properly ascribed, yet fail to account for the non-trivially different aspects of their actual expression, disappears nearly as quickly as we brought it under scrutiny. With these tools in place we can ascribe some VB to an agent using Sellarsian block quotes such that we classify their expression in terms of its functional role. If the agent uttered something non-trivially different from what is ascribed using the block quotes, then the ascription could not properly be described as correct, thus we would need to revise the ascription.

It appears as though we now have a sufficiently powerful formulation of what it would mean to give an object a functional classification. With no apparent objections ready at hand, it is now time to move on to describing how, if at all, we might be able to represent a functional classification. When we last left the discussion we were presented with two different ways that we might understand the representation of a functional classification. The first was the idea that the we would need a notion of a representation that was capable of straddling mind and world. The second, and obviously more extreme, possibility is to eliminate the notion of a representation altogether. Intuitively it seems more plausible to go

with the earlier approach since it is more modest in its scope, but before we do so, let us see if we can temper the latter.

When talking about the elimination of some theoretical tool, we can talk about an elimination as being either an ontological elimination or a methodological elimination. The first, and most extreme, is an all out ontological eliminativism. We can see a shining example of this form of eliminativism in Paul Churchland's essay "Eliminative Materialism and Propositional Attitudes." In this essay Churchland argues that the idea of propositional attitudes is so fundamentally confused that we ought to eliminate it from our ontology altogether and start anew (Churchland 1981).⁶ He sees propositional attitudes, at least as traditionally construed, to be on at the level of phlogiston theory, a theory that is not just false, but radically false, and as such it needs to be displaced, for there is no hope of reduction. The second kind of eliminativism is a methodological eliminativism. This means that we eliminate the notion from our ontology, or at least set it aside, for methodological reasons. In other words, we recognize that our current methodology precludes us from offering an adequate account of such a notion, so we set that notion aside to be the concern of another methodology. Obviously, this approach commits us to a kind of methodological pluralism. While the idea of methodological pluralism isn't obviously absurd, it would definitely be nice if we have an approach that could explain much of what we already know about cognition and at the same time be able to explain meaning, if it is, in fact, a matter of functional classification. As it turns out, DST just may be the tool that we need.

Dynamical systems theory is a hypothesis about cognition that suggests that cognition, or at least a great deal of it, functions more like a dynamical system, rather than as a computational system. A dynamical system, roughly construed, is any system with numerical states that evolve over time according to some rule (Port & van Gelder 1995). There are a number of features exhibited by dynamical systems but the most important for our purposes is that dynamical systems develop in real time. Secondly, the various aspects of a dynamical system belong, in some sense, to a collective whole. This means that the way that the various aspects develop needs to depend on one another. If any further aspects that relate to this whole in much the same way as the aspects that already belong to it, they too must be considered to actually be a part of the whole (Port

⁶ Dan Dennett likens this to flying a kite where the kite gets entangled in a tree. The kite string is so thoroughly entangled that it couldn't possibly be untangled and freed. Dennett's suggestion is that we simply cut the kite string, and attempt to fly a new kite somewhere away from the trees.

& van Gelder 1995). What does this have to do with functional classification? As we mentioned before, for a system to exhibit, or represent, a functional classification it must, to some degree, straddle mind and world. In the context of DST this means that both mind and world belong to the same system since the interaction of the agent with the world is necessarily required for the recognition of a functional classification, as in the case of a Language Entry Transition. Furthermore, if we intend to describe meaning in such a framework, we might expect the meaning system to exhibit some features of other dynamical systems. As it turns out, we do. For example, one key feature to many dynamical systems is something Clark has termed “continuous reciprocal causation” (CRC) (Clark 2001). Continuous reciprocal causation is exhibited by, among others, systems such as the Watt governor (van Gelder 1995).

The Watt governor is a system devised to govern, that is, keep constant, the speed of a flywheel attached to a steam engine. The device is composed of a vertical spindle geared into a flywheel with two arms attached to the spindle. Attached to the arms, there are two weights. As the rate of spin in the spindle increases, these two weights are driven upward due to centrifugal force. As the weights raise, they pinch a valve so that the engine reduces the amount of steam output. As a result, the weights fall to a lower position, releasing pressure on the valve and increasing the rate of steam output, which causes the cycle to continue. Within this system the weights and valve are in a CRC relationship. This means that as the valve causes more output, it causes the weights to rise, which causes the valve to close, which causes to output to decrease, which causes the weights to fall, which causes to valve to open and the output to increase and so on. Another noteworthy feature of such systems is the importance of time. This is particularly important because we see CRC taking place in real time, as a direct function of each of the involve components. If this strict time relation did not hold, neither would the system and we would no longer consider it to be a dynamical system. This is not in and of itself bad, but we would have a broken system regardless.

The main idea is that these components are involved in a coupled relationship. Analogically, we can see the notion of functional classification involving a similar kind of relationship between the environment and the agent. For example, lets say that an agent begins to participate in the patterned-governed linguistic behavior of a Language Entry Transition. The first thing that happens is an agent makes a sense impression inference from some perceptual situation. Let us now assume that this agent has made a correct inference from such a situation. When the agent expresses their inference, they will

be entitled to it by their linguistic community. The agent has now successfully made an inference and could be said to recognize the functional role of such an utterance in their conceptual economy. Now, let us suppose that some further agent made a sense impression inference from their perceptual situation. This second agent then submits his inference for the approval of his community, of which the prior agent is a member. We are now beginning to see continuous reciprocal causation. In this case, the relationship involves the linguistic community and the agents that compose that community. The agent leans on their community's appraisal and in so doing further informs later appraisals in that community. Again, as with the Watt governor, we see time playing a distinctive role in this process. We would not consider the latter agent to actually recognize the functional classification of their utterance if they were presented with some perceptual situation, but then made an inference two weeks later. At the very least we would have reason to doubt this agent's cognitive capacities and/or proficiency in the language.

As we shift to the framework of dynamical systems we can begin to see how a number of Sellars's insight could potentially be cashed out in such a system. We have noted that Sellars's commitment to inner conceptual episodes cannot quite capture the insights provided by understanding meaning as functional classification. Luckily, Sellars was able to provide us with an account that doesn't necessarily need to make use of inner conceptual episodes. Sellars doesn't seem able to cash out his insights in terms of inner conceptual episodes because some of the components that belong to functional classification, such as Language Entry Transitions, appear to not merely be inner conceptual episodes, since they require some way of accounting for the interaction with the environment. As a remedy, I have suggested understanding this type of functional classification as analogically related to a dynamical system. Viewed as such, we see that we have the tools to understand what it would mean for meaning to be functional classification. Taking our lesson from Sellars, we have also provided a framework that allows us to be relatively agnostic about the form of the representations, at least in the sense that we have a framework that could allow us to either make sense of a representations that straddles mind and world, or eliminate the notion of representations altogether. This is possible by talking about representation, and the manipulations of representations, as a fine-grained theoretical framework for describing VB. The difference lies in ontological status that we give these theoretical posits. If we consider them to be ontologically real, then we would fall closer to the view that claims that an agent

participates in the representation of a functional classification, yet is not sufficient for such a representation. If we take something like an instrumentalist approach to these theoretical posits, then we can refer to them merely as theoretical posits, posits with no real ontological status, but useful nonetheless. Finally, as promised, we see the development of something like a symbiotic relationship between the Sellarsian tradition. The Sellarsian tradition receives a framework with which it can make sense of a functional classification, while the DST receives a sophisticated formalization of what a representation ought to be. Taken together we have a powerful tool set that will allow us to tackle many of longstanding problems within cognitive science and philosophy of mind and language. It is with these tools that I leave you to address problems within your own area of specialization.

REFERENCES

- Churchland, P.M. (1979). *Scientific Realism and the Plasticity of Mind*. Cambridge: Cambridge University Press.
- Churchland, P.M. (1981). "Eliminative Materialism and the Propositional Attitudes." *Journal of Philosophy* 78: 67-90.
- Clark, A. (1997). *Being There: Putting Brain, Body, and World Together Again*. Boston, MA: The MIT Press.
- Clark, A. (2001). *Mindware: An Introduction to the Philosophy of Cognitive Science*. Oxford: Oxford University Press.
- Dennett, D.C. (1991). *Consciousness Explained*. Boston, MA: Back Bay Books.
- Hofstadter, D.R. (1999). *Godel, Escher, Bach: an Eternal Golden Braid*. New York, NY: Basic Books, Inc.
- Kelso, J.A.S. (1995). *Dynamic Patterns: The Self-organization of Brain and Behavior*. Boston, MA: The MIT Press.
- Port, R.F., van Gelder, T. (1995). *Mind as Motion: Explorations in the Dynamics of Cognition*. Boston, MA: The MIT Press.
- Resnick, M. (1994). *Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds*. Boston, MA: The MIT Press.
- Sellars, W. (1974). "Meaning as Functional Classification: A Perspective on the Relation of Syntax to Semantics." *Synthese* 27: 417-437.
- Sellars, W. (1991). *Science, Perception, and Reality*. Atascadero, CA: Ridgeview Publishing Company.
- Sellars, W. (1992). *Science and Metaphysics*. Atascadero, CA: Ridgeview Publishing Company.
- Sellars, W. (1997). *Empiricism and the Philosophy of Mind* (with an introduction by Richard Rorty and a Study Guide by Robert Brandom). Cambridge, MA: Harvard University Press.