

THE ROLE OF PICTURING IN SELLARS'S PRACTICAL PHILOSOPHY

Jeremy Randel Koons
Georgetown University – Qatar
koonsj@georgetown.edu

Carl B. Sachs
Marymount University
csachs@marymount.edu

forthcoming in *Journal of Philosophical Research* **47** (2022)
Online First: <https://doi.org/10.5840/jpr2022831186>

Abstract: Picturing is a poorly understood element of Sellars's philosophical project. We diagnose the problem with picturing as follows: on the one hand, it seems that it must be connected with action in order for it to do its job. On the other hand, the representational states of a picturing system are characterized in descriptive and seemingly static terms. How can static terms be connected with action? To solve this problem, we adopt a concept from recent work in Sellarsian metaethics: the idea of a material practical inference, which (we argue) features centrally in how we picture. The key distinction is that the picturing of nonhuman animals involves only *Humean* material practical inference, in which representational states are corrected only by feedback from the environment and not from discursive interactions. The resulting view shows that Sellars's contributions to practical philosophy (especially theory of action and metaethics) cannot be separated from his contributions to philosophy of mind, language, and cognitive science. Further, the view makes it clear that picturing is neither a version of the Given, nor is it a fifth wheel to inferential role in explaining representation, but is essential to Sellars's model of how animals—including humans—represent their environment.

Keywords: Sellars, Wilfrid; picturing; inferentialism; metaethics; pragmatism

I. INTRODUCTION

Although Wilfrid Sellars was acclaimed in his own lifetime as an important philosopher, interest in his work accelerated during the Sellars renaissance that began in 1994 with Robert Brandom's

Making It Explicit and John McDowell's *Mind and World* and is growing in size and influence. (Since 2016 alone, there have been two monographs and seven anthologies devoted to Sellars's work.) Despite this renewed interest, we still seem to lack a sufficiently rich articulation of what holds together Sellars's work as a unified whole (if indeed anything does). Perhaps there is no single Ariadne's thread, in which case there is nothing amiss with philosophers mining Sellars's texts for whatever insights they find promising. If, however, that approach to Sellars predominates, there is a danger of not seeing a unity that persists amid the diversity.

We shall argue that there is a deep unity in Sellars's thought—though perhaps one that he himself was not (or not fully) aware of. This unity concerns the relation between Sellars's theoretical philosophy and his practical philosophy. Specifically, we shall argue that a correct understanding of what Sellars called “picturing,” which has been taken up by Sellarsians such as Paul Churchland and Ruth Millikan, needs to be brought into a richer conversation with Sellars's practical philosophy in two respects. First, Sellars's conception of practical inference remained formalistic. With the addition of material practical inference, based on his own work on material inference in theoretical contexts, Sellars's practical philosophy and metaethics take on a far more satisfying shape. Second, picturing is best understood as a proto-theory about the kinds of cognitive states at work in the guidance of animal behavior.

Virtually all Sellars scholarship before the publication of deVries's and O'Shea's important books (deVries 2005; O'Shea 2007) *either* emphasizes the inferentialist, anti-representationalist, and anti-Givenist strand in his thinking *or* emphasizes the role of picturing (and often associates this with Sellars's *scientia mensura* principle); and even contemporary scholarship often struggles to situate picturing within Sellars's overall philosophical framework. Our argument is that, properly understood, both of these elements—the inferentialist element and

the picturing element—are central to Sellars’s picture of rational agency in the world and that the key to their unification is to be found in Sellars’s practical philosophy. Understanding the role of picturing is critical for Sellars’s overall project, particularly for Sellarsians who embrace the inferentialist and anti-representationalist elements in his thought. These philosophers reject picturing as ill-conceived, but then struggle to find the resources within Sellars’s work to explain how the world rationally constrains our cognitive activity, how there is no mere “frictionless spinning in a void,” as McDowell repeatedly expresses the concern in *Mind and World*. Of course, in solving this puzzle, it is incumbent on us to show that Sellars’s account—properly understood—does not run afoul of the traditional objections levied against it by Sellarsians in this camp: that it is a version of the Given, that it is a fifth wheel to inferential role in explaining representation, and so on. Demonstrating the cogency of Sellarsian picturing is also essential because it helps demonstrate the continuity of human representational systems with simpler, animal representational systems (ARSs)—a project that is central to Sellars’s overall evolutionary naturalism.¹

The cogency of our argument will depend on our ability to answer the following questions:

- 1) If semantic or conceptual content is determined by inferential role, why does Sellars think you also need picturing?
- 2) Can Sellars articulate a notion of “correct picture” that doesn’t itself fall prey to the myth of the given or otherwise muddle the causal/normative or is/ought distinction?
- 3) What is the relationship between picturing as carried out by persons and picturing as carried out by creatures who are not concept users—for example, *mere* “animal representational systems,” as Sellars calls them?
- 4) If pictures contain only descriptive terms, how can they be essentially connected to action? For *purely descriptive* claims or representations cannot at the same time be *practical*, can they? Likewise,
- 5) The dominant view of ARSs by philosophers and scientists shows that those systems are essentially action oriented. But if Sellars conceives of picturing as purely descriptive,

shouldn't we conclude that his attempt to fuse picturing and ARSs was a mistake (since we cannot see how descriptive claims can be essentially action oriented)?

Given Sellars's presentation of picturing—for example, his claim that the “maps” used in picturing contain only descriptive vocabulary—it can be difficult to see how picturing is supposed to be essentially practical. We argue that the flaw in Sellars's view of picturing is largely in his presentation rather than in the theory itself. However, we concede that it is difficult to see how the purely descriptive representations involved in picturing can be essentially practical in the way required by Sellars's model. This puzzle can be overcome, however: we shall argue that this flaw in Sellars's presentation of picturing can be repaired by turning to the idea of *material practical inference*, itself indebted to Sellarsian ideas, by understanding picturing in terms of what Sellars calls “Humean inference” (MEV) and understanding Humean inference as a species of material practical inference.

II. WHAT IS SELLARSIAN PICTURING?

We shall begin with a crucial but underexplored theme in Sellars's philosophy of language: Intentionality or meaning is not a relation between mind and the world.² By this, he means that our use of “means,” “refers to,” and “is about” cannot designate a relationship between minded organisms and their ambient social and physical environments. As Sellars makes clear (e.g., in his Notre Dame lectures and in *Science and Metaphysics*), explicating the relation between mind and world by way of semantic metalanguages is precisely what leads to the epicycles of all Platonic ontology: attributes, properties, exemplifications, essences, and so on. Against the weight of the Platonic-Aristotelian tradition, Sellars insists on a radically austere nominalism according to which there are *only* concrete particulars—no real abstracta, no real universals.³ To make nominalism work, Sellars cuts off the head of the snake by denying that semantics involves

any mind-transcendent relation at all.⁴ But Sellars understands that there must be *some* mind-world relation, if only to avoid what McDowell nicely calls “frictionless spinning in a void.”

What gives friction to the mind, in Sellars’s account, is picturing.

Consider an empirical assertion such as “The red tomato is on the brown table.” First, Sellars notes that a sentence like this is “Janus faced ... as belonging to both the causal order and the order of reasons” (NAO V. §64/p. 110). As a sentence with a sense or meaning, it is inferentially articulated and as such “fraught with ‘ought’”⁵ (e.g., semantic and pragmatic norms). But we can also consider the sentence as an item belonging to the causal order, what Sellars calls a “natural linguistic object” (NLO). Picturing with respect to language is an isomorphism among two kinds of objects: the noises and marks of NLOs, and “extralinguistic objects” (ELOs). Thus, to understand how a sentence pictures, we need to consider the sentence *not* as an empirical assertion subject to norms but as an utterance, a spatiotemporal event embedded in a causal nexus that includes the habits of the animals that use that utterance in their navigation of their environment.

Since Sellars borrowed the term “picturing” from his reading of Wittgenstein’s *Tractatus* (and perhaps knew of its pre-Wittgensteinian heritage), we want to emphasize how Sellars synthesizes Wittgenstein’s idea of picturing with what the behaviorist psychologist Edward Chace Tolman called “cognitive maps” (1948). Through many experiments with maze-learning behavior in rats and the rigorous testing of his observations against the stimulus-response theory of learning, Tolman hypothesized that animals navigate their environment via an action-guiding model-like representation of that environment that continually updates the animal’s location relative to its goal and presents the animal with possible movements in response to presented stimuli. Sellars refers to cognitive maps as early as his essay “Language, Rules, and Behavior”

(1949), eleven years before he first refers to picturing in “Being and Being Known” (1960). Subsequently, Sellars uses “picturing” and “mapping” almost interchangeably. We suggest that the Rosetta stone to Sellars’s theory of picturing lies in his insight that Tolman’s theory of cognitive maps enables a naturalistically respectable understanding of Wittgenstein’s account of picturing as a *nonsemantic* language-world relation.

Cognition, then, can be understood in a *first* approximation in terms of how a map pictures the world. Yet care is needed in explicating the concept of “maps” in any theory of animal cognition. First, items on a map are complex objects that convey a wealth of information about the correlated ELO in virtue of their sign-design. Thus, to use Sellars’s example from NAO, a circle on a map represents Urbana’s population by being a certain size and represents Urbana’s distance and direction from Chicago by locating the Urbana circle in such-and-such a position relative to the circle that represents Chicago. Importantly, Sellars spends a lot of time (for example, in NAO) talking about the dispensability of predicates to argue that sentences can also be seen as complex objects. His goal here is to dispel the notion that sentences that have a subject-predicate form express a relation between two different objects (e.g., redness and a tomato), which would make an isomorphism between this relational object and an ELO harder to map or understand. (Not incidentally, it would also threaten to turn predicates into pictures of properties, which would lead to realism about universals and conflict with Sellars’s nominalism.) We could represent “is red” by writing a word in boldface and “is brown” by writing a word in italics; and we can represent “x is on top of y” by placing x above y. Thus, “The red tomato is on top of the brown table” could be expressed by the following (predicate-free) expression:

tomato
table

It is *this* object—considered as an NLO—that has (in virtue of the ought-to-be norms of language) a correlate, in the natural order, with an ELO (namely, red-tomato-on-brown-table).

Picturing is intended to solve a number of interrelated problems in Sellars's philosophical view, the foremost of which is this: if semantic content is determined *entirely* by functional roles within a language, then talk about meaning cannot secure a mind-world relation. This is not to deny that there is a mind-world relation; it is to deny that this relation is a semantic relation. Picturing is the nonsemantic mind-world relation that Sellars needs once we realize that grounding any mind-world relation in semantics is the slippery slope to realism about universals and the subsequent need for the rationalist to invoke "givenness" in accounting for our awareness of them.⁶

In the closing paragraph of "Being and Being Known," Sellars asks, "What sort of thing is the intellect as belonging to the real order?" and responds to his own question, "I submit that as belonging to the real order it is the central nervous system and that recent cybernetic theory throws light on the way in which cerebral patterns and dispositions picture the world" (BBK §59/p. 59).⁷ Whatever it means to say that we are *rational* animals, we are at least *animals*, embodied creatures in a physical environment, and so our philosophical frameworks must combine philosophical insights about semantic content with empirical discoveries about cognitive mechanisms. Hence, Sellars urges that "There is an isomorphism in the real order [i.e., picturing] between the developed intellect and the world, an isomorphism which is a necessary condition of the intellect's intentionality as signifying the real order" (BBK §31/p. 50). Just as important, Sellars's notion of empirical truth is ultimately grounded in correct picturing: "The *criterion* of the correctness of the performance of asserting a basic matter-of-factual proposition is the correctness of the proposition *qua* picture ... [T]he *correctness* of the picture is not defined

in terms of the *correctness* of a performance but vice versa” (SM V.§57/p. 136)—that is, the correctness of the performance is defined in terms of the correctness of the picture.

Thus, we can see that picturing is intended to answer the following questions:

- A. How is animal cognition integrated into the environments it represents in a way that allows for representation?
- B. To what extent can we reconstruct our manifest-image concept of ourselves as “rational animals” in terms of the theory of animal cognition that emerges from (A)?

These questions will now be addressed in terms of how picturing works to explain rational agency in the world.

III. HOW DOES PICTURING WORK?

Despite picturing’s central role in Sellars’s philosophy, many Sellarsians have concluded that it is unworkable. Even Jay Rosenberg, an early defender of picturing, abandoned the idea for much of his career before repenting of his dissent later in his life.⁸ Thus, it is incumbent on us to explain, at least briefly, how picturing is supposed to accomplish the tasks set out for it in the previous section. Doing so will serve as a bridge between our answer to the first question from the introduction—What is picturing, and why does Sellars need it?—and also the second question—Can Sellars articulate a workable notion of *correct* picturing?

Sellars insists that our empirical concepts should be answerable to the way things are. But given his inferential role semantics, the idea that thoughts with empirical content should be answerable to how things are has two distinct requirements:

- a) These internal states, considered from a descriptive point of view, systematically covary with states of the environment.

- b) These same states are caught up in an inferential “superstructure” so that they are at the same time intentional states,⁹ states with propositional content.

Thus, intentional empirical content requires the presence of states that are Janus-faced, as Sellars calls it—belonging to both the causal and the intentional order in the precise way described by (a) and (b) above. Put otherwise, while (b) determines that thoughts have *content*, (a) determines that thoughts have *empirical* content (if they do).

To understand how isomorphism is supposed to tie empirical content to the world, we will first consider—then reject—a simplistic view of what isomorphism involves. Sellars wants to say that the NLOs we produce in response to objects in the environment will come to “map” objects in the environment, in that there will be causally regular connections between ELOs $O_1, O_2 \dots O_n$ and NLOs $L_1, L_2 \dots L_n$. That is all that Sellars means by there being an “isomorphism” between objects in the real order (namely ELOs and NLOs)—simply, that O_1 s map onto L_1 s in the sense of L_1 s are caused by O_1 s; O_2 s map onto L_2 s in the sense that L_2 s are caused by O_2 s; and so on. (To be strictly accurate, by the 1968 publication of SM, Sellars regards this isomorphism as a second-order relation, one that holds between two sets of relations.)

Consider, however, the following sets of items, purely as static items, and abstracting from any causal relations that might link As to Bs:

A_1	B_1
A_2	B_2
A_3	B_3
...	...
A_n	B_n

In other words, there are relations among As, and relations among Bs, but no causal relations between As and Bs. Can we say that the relations among the As are isomorphic to the relations among Bs? Are these sets of objects isomorphic with each other? How are we to map As onto Bs? One need not have read any Wittgenstein to recognize that there are an indefinite

number of ways to map As onto Bs, and while a natural way suggests itself to the reader, that is an artifact of our representational system. If we consider the As and Bs purely as static objects, independently of causal connections linking the two sets, there is no unique mapping or isomorphism between them.

Sellars's response to this objection lies in how he thinks isomorphisms are generated through *action*. In BBK, Sellars introduces the idea of an "anthropoid robot of the future" that constructs a picture of its environment by detecting objects and their relations as it moves about in its environment. In this example, the isomorphism is realized *only* in the actual behavior of the cognitive system that correlates As and Bs through its behavior. Thus, Sellars writes that the robot's picturing has a specific embodied, dynamic character:

This picturing cannot be abstracted from the mechanical and electronic processes in which the [robot's] tape is caught up. The patterns on the tape do not picture the robot's environment merely by virtue of being patterns on the tape. In Wittgenstein's phrase, the "method of projection" of the "map" involves the manner in which the patterns on the tape are added to, scanned, and responded to by the other components of the robot. It is a map only by virtue of the physical *habitus* of the robot, i.e. by virtue of mechanical and electronic propensities which are rooted, ultimately, in its wiring diagram. A distant analogy to this picturing is the way in which the wavy groove of a phonograph record pictures the music which it can reproduce. This picturing also cannot be abstracted from the procedures involved in making and playing the record. (BBK §40/p. 53)

These considerations concerning isomorphisms raise a number of thorny questions—chief among them, "What (if anything) determines a correct picture?"—that we will return to shortly.

Although Sellars most frequently uses maps to illustrate the notion of picturing, this analogy is potentially misleading. For one thing, NLOs are isomorphic to ELOs in *four dimensions*, not merely three (much less two!). Thus, we picture objects not merely in their *spatial* relations, but also in their *temporal* relations:

Thus, the point of picturing is to place ourselves, *qua* embodied creatures, in a four-dimensional manifold so that we may navigate this manifold *as practical animals*. The “anthropoid robot of the future” from BBK and the guided missile from NAO are systems that picture for essentially *practical* purposes: the robot pictures to explore its environment, and the missile pictures to destroy its target. Representation and practical activity are inextricably intertwined: it is *as practical animals* that we represent, and the function of representations is to guide practical activity. Hence, picturing—and the type of representation it makes possible—must be understood in the context of practical activity. Picturing, while it is a representational relation, serves a *fundamentally practical purpose* in the life of the picturing animal—including the life of the picturing featherless biped. We will return to this topic—picturing and the practical—shortly. But first, we must address one of the most difficult elements of Sellars’s theory of picturing.

IV. CORRECT PICTURE?

Now we are ready to tackle the second question: Can Sellars articulate a notion of “correct picture” that doesn’t itself elide the distinction between the normative and the natural, and thus fall prey to his own criticisms of that elision?¹² A thorny issue—and perhaps the chief reason that many Sellarsians have abandoned picturing as obscure and unworkable—is the issue of *correct* picturing. Sellars is clear that picturing is a relation of items in the real order, not the conceptual order. And so it seems impossible to give an account of how any sort of relation among items in the real order could be correct or incorrect. Relations in the real order obtain or don’t obtain, but they aren’t right or wrong, correct or incorrect. Sellars puts this objection to himself in NAO:

[A]n objection can be raised to the whole enterprise. For, it might be argued, even if it were made to work, it could not do what I want it to do. For, surely, I have at best indicated how a structure of natural-linguistic objects might correspond, by virtue of certain “rules of projection” to a structure of

nonlinguistic objects. But to say that a manifold of linguistic objects *correctly* pictures a manifold of nonlinguistic objects is no longer to consider them as mere “natural-linguistic objects”...but to consider them as linguistic objects proper, and to say that they are *true*. (NAO V.§114/p. 125)

He further elaborates on the objection a few paragraphs later, writing, “[I]t is objected that to speak of a linguistic structure as a *correct* projection is to use normative language and, therefore, to violate the terms of the problem which was to define ‘picturing’ as a relation in *rerum natura*” (NAO V.§116/p. 125).

Sellars’s reply relies on what O’Shea calls the “norm-nature metaprinciple” (O’Shea 2007, p. 50): “Espousal of principles is reflected in uniformities of performance” (NAO V.§95/p. 119). He elaborates on what this means a bit later in the same chapter:

[W]hile to say of a projection that it is *correct* is, indeed, to use normative language, the principle which, it will be remembered, I am taking as axiomatic assures us that corresponding to every espoused principle of correctness there is a matter-of-factual uniformity in performance. And it is such uniformities, which link natural-linguistic objects with one another and with the objects of which they are the linguistic projections, that constitute picturing as a relation of matter of fact between objects in the natural order. (NAO V.§116/p. 125)

Here is how this is supposed to work: Consider, for example, language-entry moves, which are governed by “ought-to-be’s” like the following:

Only red objects should be responded to with “red”¹³

Trainers (NAO IV. §§24–30/pp. 68–69) condition language-learners so that they respond only to red objects with “red” by reinforcing this tendency, and extinguishing contrary tendencies (e.g., tendencies to respond to red objects with “blue”). Thus, our language learners come to correctly picture the world insofar as they come to conform to the relevant semantic rules, e.g., respond to red ELOs by producing the NLO “red.” (Of course, this isn’t *merely* an NLO; as Sellars is at pains to emphasize, the symbol “red” must be both a learned response to stimulus and also be

caught up in the web of inferential proprieties—for it to count as an NLO in the first place, and for the being producing it to count as a user of empirical concepts.)

This elaboration of “correctness of picture” is bound to breed dissatisfaction—and indeed it has. For how can it play the role it is supposed to play within Sellars’s philosophy of language? As we have already noted, Sellars’s notion of truth is ultimately grounded in correct picturing: “The *criterion* of the correctness of the performance of asserting a basic matter-of-factual proposition is the correctness of the proposition *qua* picture ... [T]he *correctness* of the picture is not defined in terms of the *correctness* of a performance but vice versa” (SM V.§57/p. 136). But we see now that correct picturing is itself defined in terms of compliance with a semantic rule—that is, is defined in terms of the rules of the language. How can picturing serve as an external constraint on language, if our only access to whether our picture is correct or not is from within language and the semantic rules of our language?¹⁴

Problems analogous to this have always presented thorny difficulties for philosophers who, like Sellars, reject givenness: all correctness seems to be confined to the circle of language; we are threatened with “frictionless spinning in a void,” to use McDowell’s metaphor. The best way to address this difficulty is head-on: yes, we have no access to reality unmediated by our conceptual scheme; but no, this doesn’t render the picturing constraint empty.

Let us begin with the first claim. Jay Rosenberg’s early articulate defense of Sellarsian picturing makes clear that picturing does not provide a set of ELOs outside of conceptual space, to which we can compare our conceptual structure in order to assess the adequacy of the latter. Whether our conceptual structure correctly pictures extralinguistic reality is something that can only be assessed from *within* this conceptual structure:

[C]orrespondence as protocorrelational isomorphism is the *outcome* of the evolution of representational systems, but cannot be the *grounds* upon which a

choice among representational systems is predicated, for that our system of representations is, in extension, protocorrelated with the world it represents can be known to be the case only, as it were, *indirectly*—through our knowing that the system of linguistic normatives giving rise to the system of natural linguistic objects protocorrelated with the world which they thereby represent is increasingly explanatorily adequate to that world as it impinges upon us through continuing experience. (Rosenberg 1974, p. 121)

We have no independent access to whether there is an isomorphism between NLOs and ELOs, except through our ongoing engagement with the world via conceptually mediated representations. Rosenberg, in revisiting the notion of picturing more than 30 years later, elaborates essentially the same view: “Since (as we are assuming) ‘a is red’ belongs to *our* representational system, its picturing, so to speak, takes care of itself. In order to determine whether ‘a is red’ is a *correct* picture ... then, the only thing that we need to *find out* is whether a is red” (Rosenberg 2007, p. 120).

Yet this doesn’t render the concept of picturing empty; picturing serves not merely as a condition on the possibility of representation, but as the empirical constraint on such representation. For example, suppose Isabella’s picture of the world contains the following (incorrect) elements:

Chicago-south-of-Urbana

Match is dipped in water at t , match is lit at $t + \Delta t$

Her initial representations and her subsequent volitions (“I am in Urbana, I am driving south”; “I am dipping this match in water”) will fail to be followed by the representations indicated on her “map” (“I am in Chicago”; “the match is lit”). Thus, the failure of her map to be isomorphic to the world—that is, the failure of her map to *correctly picture* the world—will result in a failure of her representational system to adequately represent, and it will therefore present a practical need to revise her picture in the direction of greater adequacy. Thus, it is the *incorrectness* of the

picturing relationship that will explain the pressure toward revision; revisions toward greater explanatory completeness and elimination of anomalous experiences will themselves lead us in the direction of increasing isomorphism between NLOs and ELOs. The practical need to generate isomorphisms—to have it so that (e.g.) representations on one’s map (“I am in Chicago”) correlate with actual states of the world (being in Chicago)—is what exerts pressure toward correct picturing, and this practical need itself stems from the need to be able to use one’s map to navigate one’s environment. That is what maps are for (and this is a crucial element of how picturing is supposed to help tie language to the world and guarantee empirical representational content). And so while *correctness* of picturing acts as a *regulative idea* directing us to revise our theoretical apparatus, *incorrectness* of picturing is the *concrete mechanism* driving this revision—driving it toward greater isomorphism. Crucially, it is because of the essentially practical role of picturing in the life of the animal that the consistent pressure will be toward revising incorrect maps in the direction of correctness or accuracy. As we argued earlier, the point of picturing is to place ourselves, *qua* embodied creatures, in a four-dimensional manifold so that we may navigate this manifold *as practical animals*, and so failure to revise in the direction of correct picturing will consistently inhibit practical worldly success. This is why “It is because a language is the evolutionary product of the continuing causal impact of anomalous experiential inputs to the community of representers that the structure of its extensional realization tends increasingly towards protocorrelational isomorphism with the world represented” (Rosenberg 1974, p. 119).

V. PICTURING AND THE PRACTICAL

Thus far, we have established that picturing enables embodied creatures to act in physical and social environments: “One doesn’t have to actually use [maps] to go to the places they represent in order for them to be maps, but the point of being a map is to translate into sentences which dovetail with *practical* discourse about getting from point A to point B” (NAO V.§77/p. 114). So, for example, in discussing Sellarsian picturing, Huebner notes that “Picturing is an *activity*, which depends on the use of a cognitive map to actualize one of an agent’s capacities for acting in an embodied, embedded, and situated way; and since cognitive maps provide a representation of an agent’s place in the world that *they have encountered*, they will always be deeply tied to that agent’s capacities for action” (Huebner 2018, p. 14; emphasis original).

Having defended picturing up to a point, we must now squarely face the following difficulty: Despite our stress on a relation between picturing and practical inference, Sellars explicitly and consistently holds that picturing or mapping lacks any practical or action-oriented vocabulary:

The vocabulary of a map is extremely limited, lacking, for example, logical connectives. It is equally important to note that it lacks words for actions. Thus, although a map is for use in travelling, there are no words for ‘to go forward’, ‘to turn right’, etc. Thus, even if the map tells us that Chicago is north of Urbana, it is only in the language to which we translate the map that we get

Going north-east from Urbana is going toward Chicago. If I am in Urbana and want to get to Chicago, I should first go north on Route 89.
(NAO V.§76/p. 114)

If Sellars takes cognitive systems as comprised of maplike representations, but maplike representations are purely descriptive (and do not contain any “words for action”), then it is difficult to see how, on Sellars’s view, picturing is an essentially practical activity. His account of picturing emphasizes atomic descriptive sentences and the absence of action-oriented

vocabulary. How are these maplike representations essentially practical if they are also purely descriptive? Sellars's view seems radically out of step with how picturing actually *works* in both animal minds and in discursive representers such as ourselves.

Our claim is that Sellars does intend for these maplike representations to be essentially practical. As we have been emphasizing, picturing itself is essentially practical; the point of mapping the environment is for purposes of agency. However, Sellars's own account of how, exactly, these descriptive representations can be essentially practical is obscure, which might lead one to conclude that Sellars's model of picturing is untenable. The next major step in our argument is thus to show how the maplike representations involved in picturing are essentially tied to action. We will do so by drawing upon Sellars's metaethics, or more specifically, a philosophically corrected version of Sellars's metaethics—albeit a correction that is itself indebted to Sellars.

V.A. MATERIAL PRACTICAL INFERENCE

To motivate an account that essentially connects picturing to action—and more specifically, to motivate the idea that what appears grammatically as a purely descriptive sentence can at the same time be a practical sentence, in that it entails (psychologically and perhaps also normatively) an intention or a volition—we shall need to motivate a particular view of practical inference, one that Sellars did not endorse but that he had overwhelming reason to endorse, given his other philosophical commitments.¹⁵

Consider the following inference:

P: Fido is a dog.

Q: Therefore, Fido is a mammal.

Readers of Sellars and Brandom will of course be familiar with the debate over whether to interpret such inferences as *enthymemes*—that is, as relying on a suppressed premise of the form:

$P \rightarrow Q$: If Fido is a dog, then Fido is a mammal.

We will call philosophers who do read such inferences as enthymematic *formalists*, as they think that only deductively valid inferences are legitimate.

Sellars and Brandom, of course, reject this line of reasoning, citing a line of argumentation that goes all the way back to Lewis Carroll (1895).¹⁶ Suppose you are presented with the following argument:

P: Fido is a dog.

Q: Therefore, Fido is a mammal.

The formalist argues that you cannot infer a conclusion from a premise without an explicit conditional linking premise and conclusion. Thus,

$P, P \rightarrow Q$; therefore, Q.

But now, argues Carroll, we have merely pushed the problem back a step. We have two premises we need to link to the conclusion via a conditional. No problem! We'll just insert a new conditional:

$(P \ \& \ [P \rightarrow Q]) \rightarrow Q$

So our new argument reads as follows:

$P, P \rightarrow Q, (P \ \& \ [P \rightarrow Q]) \rightarrow Q$; therefore, Q.

Of course, now we have three premises that need to be linked to the conclusion, and the problem is not solved. The formalist demand threatens all reasoning with an infinite regress.

As Brandom (2000) notes, we can construct formally valid inferences out of materially valid inferences by taking formally valid inferences as making explicit inferential proprieties that

are implicit in practice. But, as Carroll has shown, there is no converse route that constructs materially valid inferences out of formally valid inferences. Hence, we must take materially inferential goodness as primitive and understand formal inferential goodness as *expressive* of inferential proprieties that are implicit in these materially good inferences.

Brandom urges that we adopt this route—taking material goodness as primitive—in both theoretical and practical reasoning. Thus, consider the following moral principle, which we will state in the Sellarsian language of we-intentions:

Shall_{we} (I do A_i in C_i)

If we consistently read Sellars as an expressivist, and not a formalist, then we should understand the goodness of this material inferential principle as parasitic on the goodness of following material inferential propriety:

I am in C_i
Therefore, Shall_{we} (I do A_i)

Unfortunately, this is not how Sellars reads the above moral principle. Instead, Sellars reads the above inference as an enthymeme, reliant on the above moral principle, which is itself derived from the supreme principle of morality—namely, “Shall_{we} (I promote the welfare of each and every one of us, all relevant things considered).”

One of us has argued at length (Koons 2019, chapter 14) that a consistent Sellarsian should reject formalism in practical reasoning for the exact same reasons that Sellars himself rejects formalism in theoretical reasoning. Thus, the practical inference from the previous paragraph should be understood as a *materially valid practical inference*, not an enthymeme, and in no need of supplementation by a conditional in order to be a good practical inference.

This has consequences for how we understand moral perception in Sellarsian ethical theory. Thus, consider the following example (borrowed from Koons 2019, p. 244). Suppose our

earlier practical reasoner, Isabella, is confronted with the following scenario, with a list of possibly salient facts (and reasons for their salience):

There is a cool lake (and I am thirsty)
There is a glorious view (and I have a camera)
There is a shady tree (and it is hot)
There is a beautiful patch of wild orchids (and I am an amateur botanist)
There is an angry hippopotamus charging toward me, with obvious intent to wound and trample

In forming her conception of the situation (and of what she ought to do in this situation), Isabella notices a particular fact that is salient (and she may regard other facts as not salient at all—indeed, she may not even *notice* other facts about the situation, facts that in different circumstances might be salient, or even conclusively reason-giving). Thus, when Isabella surveys the situation, she takes the overridingly salient fact of the situation to be that there is an angry hippopotamus charging toward her, with obvious intent to maim and kill. But the fact that she perceives this fact as salient cannot be understood apart from her overarching concerns—her concern for her own life and limb, for example. Thus, Isabella’s very conception of the situation—how she perceives the situation, what she sees as salient—is conditioned by her concerns and commitments, and it cannot be understood apart from them. Furthermore, we understand this feature of the situation as salient precisely because it is connected to her concerns about preserving her life and limb, and therefore to action—say, immediately escaping danger by leaping across a moat, so as not to be wounded and trampled. The salience emerges from engagement with Isabella’s particular concerns and demands a particular action. Since cognition itself is already laden with evaluative and practical concepts and concerns, a practical inference such as

I am in C_i
Therefore, Shall_{we} (I do A_i)

is materially valid. Rejecting formalism means understanding cognition as practical in a very specific way: It means understanding our conception of the world as engaging our concerns and values.¹⁷

The upshot of this discussion is that a representation can be essentially practical without containing any “action words.” Isabella’s representation that there is an angry hippopotamus charging toward her, with obvious intent to wound and trample is in a sense purely descriptive—if we abstract away from the ways in which her cognition of this involves her interests and concerns. But *qua* finite agent, this cognition is imbued with such concerns, and hence it is an essentially practical cognition, albeit in a sense a descriptive one.

Hopefully, now, the view is beginning to emerge of how maplike states can be descriptive representations but can also (given the various concerns or background intentions of the organism) be action guiding, although the full picture will only emerge in the second half of this essay. With this conception of the fundamentally practical nature of human cognition in the background, we will now show how to understand this conception in light of our earlier articulation of picturing. In doing so, we will demonstrate the continuity of human and animal representational systems and demonstrate how human representational systems are integrated into the causal order. Second, we will demonstrate how systems (such as animal cognition) that do not traffic in discursively articulable concepts can nevertheless be *representational* systems—a thorny problem for some inferentialists. Finally—and crucially—this conception will allow us to demonstrate how such representational systems are essentially practical, and how picturing and representation are essentially tied to action both in ARSs and in discursively structured representational systems of the sort that constitute us as persons. Thus, let us now turn to

answering our third question: What is the relationship between picturing as carried out by persons and picturing as carried out by creatures who are not concept users?

V.B. FROM RDR TO ARS

The first step in our reconstruction of picturing as practical cognition is to adopt Sellars's concept of ARSs (MEV) for understanding animal cognition. We need to further develop Sellars's account of ARSs as essentially action-guiding, value-laden picturing systems. To do so, we start with the basic fact that a picturing system must be able to generate states that covary with states of its environment. We will follow Brandom in calling such state-generators *reliable differential responders* (RDRs). An RDR is just that—it will reliably respond to objects with “reports.” (We use scare quotes because being an RDR is not sufficient to count as a representational system.) In this sense, almost anything can serve in some dimension as an RDR—a parrot can be an RDR, and so can a thermometer, or even a chunk of iron rusting in water. A system considered purely as an RDR is static, however—all it does is respond to objects by producing a particular state. This bears a superficial relation to picturing—after all, an RDR will respond to object A_1 by producing state B_1 , to object A_2 by producing state B_2 , and so on. But picturing is a form of representation, and Sellars is clear that for a system to be a *representational* system, it has to be more than a mere RDR: “to be a representational state, a state of an organism must be the manifestation of a system of dispositions and propensities by virtue of which the organism constructs *maps* of itself in its environment, and locates itself and its behavior on the map” (MEV §56/p. 336).

To transform an RDR into an ARS, the RDR must acquire the capacity to engage in “proto-inferences,” what Sellars calls “Humean inference” as distinct from the “Aristotelian

inference” of genuinely rational animals. The distinction between the two kinds of inference corresponds to a difference between logic-using and non-logic-using representational systems (MEV §§78ff/pp. 340ff). A non-logic-using representational system can engage in association of ideas represented as follows:

‘Smoke here’, ‘fire nearby’

On the other hand, a logic-using representational system can engage in Aristotelian inference, which involves the use of logical vocabulary, and hence the ability to generalize and to access a metalanguage. So, for example, a logic-using representational system can infer using quantified premises involving logical vocabulary; e.g.,

“If smoke anywhere, then fire nearby there”

The crucial point here is that to represent something is to represent it *as* being *somehow*—something that mere differential responders do not do.¹⁸ But on Sellars’s inferentialism about semantic content, such content can only come from the material inferential proprieties governing the use of an item, and so for a system to be a representational system, the states produced by the RDR must at least be caught up in a web of Humean inferences connecting these states to other internal states of the organism and to output (i.e., action) states. It is only by engaging in this complex web of entry, intra, and exit transitions that an organism can be treated as *representing*—and in particular as representing its environment and its relation to that environment.

This, then, is the first step in Sellars’s project of situating the mind in nature. He demonstrates how ARSs are built out of RDRs: ARSs have RDRs as an essential component. But let us try to apply here some of the lessons we learned in the previous section (V.A.). Implicit there was a rejection of what Marc Lange (2000) calls the “layer-cake” view of practical

cognition: a purely factual/descriptive cognition of the world with an evaluative or conative conception added to it. We agree with McDowell, Lange, and other Sellarsians on this point with regard to human practical cognition, but this recognition should be built into the very foundations of our account of how ARSs work. In this, we are following McDowell, who takes Sellars’s argument for the *conceptual* saturation of perception and applies it to the *axiological* saturation of perception.¹⁹

How do we resist the “layer-cake” view when interpreting ARSs? As we have already noted, on the Kantian view of cognition, the acquisition of new conceptual capacities can transform all of the upstream cognitive capacities.²⁰ Thus, a *mere* reliable differential response can *become* a conceptually contentful observation report because one acquires a battery of inferential abilities. And this provides a model for how we should see ARSs as transforming the outputs of RDRs. In virtue of acquiring a battery of Humean inferential capacities—both the capacity to associate one representation with another and also the ability to transition from representation to volition—the initial reliable differential responses are transformed from mere static elements into elements laden with the possibility of action. Thus, consider again a reliable differential responder with internal states isomorphic with objects in the world:

S ₁	O ₁
S ₂	O ₂
S ₃	O ₃
...	...
S _n	O _n

Considered purely as the outputs of an RDR, the internal states S₁...S_n don’t represent O₁...O_n *as* being any way; they are merely correlational, not genuinely *representational*. But if we consider the creature as an ARS, its internal states are not only generated by objects atomistically but also connected among themselves via series of Humean inferences, as well as being connected to

representation-entry transitions (perceptions) and representation-exit transitions (volitions).²¹ Suppose an animal is characterized in terms of the simultaneity of internal states S_1 and S_2 (say, hunger and cognition of a certain landmark) and thereby produces a volition (to navigate to a certain bush, which it represents with another internal state), and then searches for certain objects (berries) and produces another volition (to eat these objects). Thus, the creature represents the landmark and the bush *as* being in this spatial relation (or navigable in this way); we can say the creature represents the berries *as* being accessible from this landmark and *as* edible, and so on. As we see it, the initial states S_1 and S_2 can only be characterized as having the representational purport we attributed to them in virtue of their involvement in these and other proto-inferential moves.

A significant advantage of the emerging Sellarsian theory is that it offers a plausible account of animal representation, whereas more “intellectualist” inferentialists like Brandom often have more difficulty doing so, given that animals don’t have Brandomian commitments and entitlements. We concede to the Brandomian that nonlinguistic animals make only Humean inferences; they cannot represent these inferential proprieties in a metalanguage; and so on. But this element of Sellars’s account shores up a potential weakness with inferentialist theories: it allows us to attribute *mental states*—intentional, representational states—to non-concept-using animals. We think this is a condition of adequacy for any account of the mental—a condition that not all accounts obviously satisfy.

At this point, it is necessary that we tip our hand as to where Sellars is going with this. A chief motivation for Sellars’s discussion is to demonstrate how human cognitive activity is largely continuous with, and emerges out of, cognitive processes that exist in nonhuman animals. We, too, are ARSs—although we are ARSs with unique capacities. We might say that nonhuman

animals are *associative* ARSs—henceforth, ARS_{AS}—whereas humans are *discursive* ARSs, or ARS_{DS}.

These elements—the uptake of an RDR into an ARS via Humean inference and the connection to action—allow us to demonstrate how there is, in ARS_{AS}, an analog to material practical inference. Thus, we can attribute to animals transitions of the following sort:

Cognition [That there is a predator within flight distance]
(Therefore) Volition [Fleeing]

Cognition [Edible berries there]
(Therefore) Volition [Eat the berries]

Because ARS_{AS} engage in only Humean associative inferences, they do not represent these inferences to themselves, nor do their cognitions involve any logical vocabulary. But just as with human animals, ARS_{AS} should be seen as transitioning (in many cases) from cognition of a state of affairs directly to an intention or a volition. Indeed, this is crucial—for material practical inference, as deployed by human animals such as ourselves, is built out of the cognitive tools that are already present in ARS_{AS}, and so there must be something analogous to material practical inference already present in the ARS_{AS} that underlie and partially constitute concept-mongering systems such as our own. In ARS_{AS}, the connection of certain internal states to action is partially what gives them the representational content they have: the animal analogue of material practical inference. Thus, it is because the prairie dog's distinctive call is in response to an approaching predator and induces prairie dogs within hearing to seek cover that the call represents the presence of an approaching intruder and represents it *as* a threat. Formalism is even less plausible with respect to ARS_{AS} than it is with respect to human practical cognition. It is much more plausible to understand a prairie dog's cognition of the warning call as *inherently action guiding*. More generally, approaches to animal cognition focused on affordances similarly understand

animal cognition often to involve inherently practical representations—and hence to involve an analog of material practical inference, a cognition that moves directly to volition without requiring an intervening inferential premise (even assuming the inference in question to be Humean).

Thus, if the reader was worried that on our account of material practical inference (developed in section V.A.), such inferences essentially involve language, this discussion should hopefully lay such concerns to rest. The material practical inferences performed by ARS_{DS} are in general enabled by a linguistic conceptual system. But the basic cognitive hardware is the ability to cognize the world in a way that embodies affect/interest/concern and to directly transition to an intention or a volition from this cognition. Such hardware is present in ARS_{AS} as surely as it is present in ARS_{DS}—the difference being that in the former case, the material practical inferences involved are always Humean inferences.

As we have previously noted, Sellars describes observation reports as Janus-faced: as part of the natural order (i.e., as NLOs) and as part of the normative order (i.e., as language-entry transitions). Understanding the place of the human mind in nature depends on understanding these items as describable in both natural and normative terms even though the normative cannot be reduced to the natural. Just as seeing something as an NLO abstracts from its conceptual character and involves seeing it purely as an item in the natural order, seeing the internal states of an ARS_A in abstraction from Humean inferential relations allows us to consider these internal states (and their relations) only as *correlated* with analogous items and relations external to the animal. Construed solely as elements in a picturing relation, these internal states can be viewed as purely descriptive—they are (when viewed at this level of abstraction) not action guiding; they do not essentially involve any action vocabulary. Thus, considered purely as an element

isomorphic with an object in the world, the internal state of our creature doesn't represent the lion *as* dangerous. What it represents is brown-lion-crouching-in-tall-grass. However, when we consider the state as part of a system of animal representations tied to Humean associative inferences, including practical inferences, the state represents dangerous-brown-lion-crouching-in-tall-grass—and in virtue of this practical element, it is potentially connected to volition (i.e., fleeing). In other words, the puzzle with which we began—how can purely descriptive maplike representations be essentially practical?—can be put aside once we realize that the descriptivism in Sellarsian picturing is an artifact of abstracting cognitive states from their functional integration into an ARS.

At this stage, it is important to see picturing as a component of a prelogical ARS_A, where picturing consists in changes to the map in response to changes to the body/environment of the creature. The map gets updated, and every updating is available for new Humean inferences. Thus, when the rabbit acts on a volition, which we would express by attributing to it:

Shall_I *now* (travel to landmark L)

the rabbit will (in response to a perception of its new location) update its internal map by updating its relative position on the map, and this in turn updates the inferences it will make based on its instrumental goals (i.e., if confronted with danger, shelter is reachable in *this way* from L, so it will form the volition to go *this way*; but if it is hungry and wants to forage for berries, the berry bush is *that way* from L, so it will form the volition to go *that way*). These are Humean inferences, as they are volitional responses (e.g., the volition to go *that way*) to map information (being at location L) and bodily/environmental imperatives (hunger, danger, etc.). Thus, though ARS_{AS} are not logic users, they are practical animals who deploy a kind of

rationality, though it is not the kind of full-blown rational agency possessed by agents who make Aristotelian inferences.

V.C. FROM ARS_{AS} TO ARS_{DS}

On this basis, the transition from ARS_{AS} to ARS_{DS} should be relatively straightforward. Non-logic-using representational systems can make only Humean inferences—basically, associations of ideas that do not involve logical vocabulary. As Michael Tomasello writes about organisms with individual intentionality—basically, ARS_{AS}:

[T]he combinatorial processes [used in animal inferences] will include ...“logical” operations such as the conditional, “negation,” exclusion, and the like. [However, t]hese logical operations are not themselves imagistic cognitive representations but, rather, cognitive procedures (enactive, in Bruner’s terms, or operative, in Piaget’s terms) that the organism accesses only through actual use. (Tomasello 2014, p. 13)

ARS_{AS} can act in conformity with rules, but they cannot act according to their conception of a rule—a recognizably Kantian idea. A logic-using representational system, by contrast, can make Aristotelian inferences, which involve premises that make use of tools (unavailable to ARS_{AS}) such as quantifiers and logical vocabulary; e.g.,

“If smoke anywhere, then fire nearby there”

For Sellars, such claims are metalinguistic, in that they assert an inferential license between two claims. So what sets ARS_{DS} apart from ARS_{AS} is that the former can take the inferential transitions *themselves* as objects of representing, and in consequence of this they can debate (through the game of giving and asking for reasons) whether we should be bound by these—or by different—proprieties.

The capacity to access the metalanguage and to act on explicit representations of rules—as well as, crucially, the corresponding ability to subject these normative proprieties to scrutiny

and, if necessary, revision—is what makes an ARS_D out of an ARS_A. And just as with the transition from RDRs to ARS_{AS}, the transition from ARS_{AS} to ARS_{DS} transforms all of the upstream cognitive capacities. What was a mere association of ideas now becomes an inference—not because there is a different psychological process involved (there need not be) but because the ARS_D is *accountable* to others (and itself) for the inferential move from P to Q. That is, if (appropriately²²) challenged by another rational agent, the ARS_D must either defend the propriety of the inference or withdraw the inference. ARS_{AS} are not capable of such moves and not normatively accountable for their (Humean) inferential transitions. The addition of new cognitive capacities allows us to consider inferential transitions as belonging to the normative order (as well as the causal order), whereas without these additional capacities, these transitions would belong only to the causal order. The transformation of upstream cognitive capacities does not necessarily mean that these capacities are changed in their intrinsic qualities. Sometimes it is merely their normative status that alters with the acquisition of new normative capacities (e.g., the capacity to access the metalanguage).

It must be emphasized that ARS_{DS} are still chiefly makers of Humean associative inferences—and so their cognitive activities are largely continuous with those of ARS_{AS}. Sellars is clear—for example, in SRLG and NAO—that the various kinds of linguistic moves (such as language-language transitions) made by ARS_{DS} are acts, not actions, and *considered purely as transitions* they are not *intrinsically* different from the transitions made by ARS_{AS}. This is the sense in which ARS_{DS} *are* ARS_{AS}—they are vastly complicated systems of causal associations. The associations occurring within an ARS_D do not possess the kind of inferentially articulated conceptual content we deploy because of anything *intrinsic* to the transitions themselves. Instead, these transitions have such content because we can represent the transitions themselves

metalinguistically. An ARS_D largely *is* an ARS_A—albeit an ARS_A that has acquired the capacity to represent its own inferences in the medium of the metalanguage. Thus, both ARS_{AS} and ARS_{DS} can make the Humean inference, “Smoke there; therefore, fire there.” But on the inferentialist picture defended by Sellars and later by Sellarsians like Brandom, this Humean inference is conceptually contentful for ARS_{DS} because they can represent such transitions as governed by a rule (“Wherever there is smoke, there is fire”); they can use these Aristotelian inferences in the game of giving and asking for reasons (“There is smoke over there, therefore there is fire”); and they can use these metalinguistic resources to revise the propensities both to make various Humean inferences and to revise the commitments they have based on prior Humean inferences (“That’s not smoke, that’s steam”—“Oh, so probably there is no fire over there after all”). As Brandom writes, “Being subject to rules is not special to us as discursive, that is concept-applying, subjects of judgment and action. What is distinctive about us as normative creatures is the *way* in which we are subject to norms (for Kant, in the form of rules). As natural beings, we act according to rules. As rational beings, we act according to our *conceptions* of rules” (Brandom 1994, p. 30). We have in common with ARS_{AS} that we act on habits, learned associations; however, we go beyond ARS_{AS} in that we can represent the proprieties implicit in these associations in the form of explicit rules. (And in doing so, these associations come to embody *proprieties* instead of just *regularities*—this is the birth of normativity.)

The role of intersubjective accountability does not affect the intrinsic character of representational states, but it does affect how they can be held and what we can do with them. First, the transitions made by ARS_{AS} are mere *learned associations*, while ARS_{DS} make learned associations for which they are normatively *accountable*.²³ Thus, a bird learns through

behavioral conditioning that certain kinds of butterflies (e.g., monarchs and viceroys) are foul-tasting, and hence it avoids eating them. It would be incorrect to say, however, that the bird ought not eat monarch butterflies, or ought not represent the monarch as edible. To be sure, the bird will face consequences for doing so, and it may learn from those consequences (or die from them). But it will not have done something wrong in the sense of having transgressed a norm; it has not acted in a way such that other birds would chastise it for having failed to eat properly.²⁴

Thus, whereas animals receive a kind of error correction via sensorimotor feedback from their environments, only ARS_{DS} hold each other accountable for their inferences and actions. If you habitually make transitions of the sort, “x is a fish; therefore, x is a mammal,” then you are wrong according to the classificatory framework shared by other persons. You are accountable not merely to the world (which this framework pictures), but to the other persons whose epistemic-cum-semantic practice implements this picture. They may rebuke and correct you for your faulty transitions in a way that is independent of “mere” worldly consequences (such as experiencing a foul taste from eating the “wrong” insect).²⁵

Put otherwise, ARS_{AS} have only their own individual experience for revising their stock of transitions. The way a bird goes from representing a monarch *as* edible to representing it *as* inedible is by eating a monarch. (This is because animals are chiefly accountable to the world, and their accountability to each other is chiefly via worldly causal consequences, such as being poisoned or eaten.) But ARS_{DS}—who are also accountable to *each other*—have another way of revising their transitional practice: As Aristotelian reasoners, they can represent these transitions in the metalanguage and hence debate the accuracy and propriety of these transitions. Thus, we can *tell* you that “x is a fish; therefore, x is a mammal” is wrong and *argue* that you are wrong. *This is a fundamentally interpersonal and normative endeavor*, and it is an activity that is in

principle unavailable to ARS_{AS}, who lack the metalanguage that would make this activity possible.²⁶

To elaborate a bit on this difference, we will note that there is a sense in which many ARS_{AS} can “represent” their own inferences. As Tomasello notes, many animals can “go through the kind of inferential simulation that Piaget (1952) called ‘mental trial and error’: the organism imagines a potential action and its consequences” (Tomasello 2014, p. 13). For example, “a chimpanzee might simulate imaginatively what would happen if she forcefully tugged at the stick, without actually doing it” (Tomasello 2014, p. 13), or a squirrel might mentally model jumping to a distant branch (before deciding it is too far). As Tomasello notes, imagining sequences of actions and making “corrections” before actually undertaking any action is what “Dennett [1995]²⁷ calls...Popperian learning because failure means that my hypothesis ‘dies,’ not me” (Tomasello 2014, p. 14).

What distinguishes ARS_{DS} from ARS_{AS}, though, is the specific *way* in which the latter are able to represent their own inferential transitions. As Tomasello suggests, ARS_{AS} make logical transitions, but they don’t have mental representations corresponding specifically to logical vocabulary such as negation, conditionals, and so forth. ARS_{DS}, by contrast, do—this is what it *is* to be an agent who makes Aristotelian inferences. The presence of logical vocabulary allows ARS_{DS} a very specific mode of representing their inferential transitions—the *metalinguistic mode*. We cannot do justice here to the importance of this difference, but it ties in essentially with Sellars’s emphasis on the intersubjective nature of rational agency (emphasized at the end of PSIM and elsewhere): as one of us has argued elsewhere (Koons 2021), there can be no interpersonal rational resolution of disputes without the existence of a Sellarsian metalanguage (and its associated logical resources) that allows the explicit formulation of inferential relations among

claims. Thus, the creation of a rational community very much depends on this essential difference between ARS_{AS} and ARS_{DS} .

Finally, language users *qua* language users are able to make many transitions that non-language users cannot. As Brandom (1979), for example, has convincingly argued, there are many thoughts one simply cannot have if one lacks a suitable language; and there will be many transitions one simply cannot make if one is not a language user—for example, “x is an electron, so x has a negative charge.” So there is a final sense in which *some of* the various transitions can differ intrinsically between ARS_{DS} and mere ARS_{AS} .

None of this should obscure, however, the similarities between ARS_A and ARS_D inferences. For Sellars, both are causal learned associations. For Wittgensteinian reasons (see SRLG), such transitions cannot in the base case be examples of rule-following; the ability of ARS_{DS} to make Aristotelian inferences doesn’t mean that ARS_D transitions are standardly cases of rule-following. No; human behavior (including language-entry, language-exit, and language-language moves [i.e., inferrings]) largely consists, for Sellars, of *acts*, not *actions*. What is distinctive of human cognition is that although it is still constrained by the requirements of sensorimotor behavior (“base”), it is *also* characterized by norm-governed symbolic activity (“superstructure”). What distinguishes the transitions made by ARS_{DS} from the transitions made by ARS_{AS} is largely extrinsic to the transition itself: it is the presence of the metalanguage; it is the incorporation of the ARS_D into the community of persons; it is the normative accountability of the ARS_D for the transitions it makes. This isn’t something added on top of the RDR—à la the layer-cake view—but is a way in which possession of additional cognitive tools (such as a metalanguage) can radically transform “simpler” cognitive tools (such as RDR outputs and

Humean inferences) along various dimensions, such as transforming Humean associations from mere associations into normative proprieties.

VI. PUTTING IT ALL TOGETHER

We are finally in a position to answer the last two questions/puzzles we presented in the introduction to this essay—puzzles that seemed to stand in the way of understanding Sellars's theory of picturing essentially as connection to action and the practical. The puzzles were:

- 4) If pictures contain only static descriptive terms, how can they be essentially connected to action? For *purely descriptive* claims or representations cannot at the same time be *practical*, can they? Likewise,
- 5) The dominant view of ARSs by philosophers and scientists shows that ARSs are essentially action oriented. But if Sellars conceives of picturing as purely descriptive, shouldn't we conclude that his attempt to fuse picturing and ARSs was a mistake (since we cannot see how descriptive claims can be essentially action oriented)?

We now have the resources to solve both of these puzzles. Let us begin with (5). Sellars can perfectly well explain how ARSs are intrinsically action oriented. In fact, an ARS (whether it is an ARS_A or an ARS_D) is only an ARS *in virtue of being action oriented*. As we have argued, an RDR is not in itself a representational system: it is correlational, not representational. For a system to represent objects *as* being *somehow*, it has to systematically correlate these objects not only with states; it has to correlate these states with each other and with action in some systematic way. Thus, as in the example we offered earlier, it is only by using landmarks to navigate to the berry bush and eating berries that the animal represents the berries *as* being edible, the bush and the landmarks *as* being in this spatial relationship, and so on. So ARSs are essentially practical, because they wouldn't be *representational* if they weren't practical.

Having addressed (5), we can move on to (4). Recall that for Sellars, the internal states of the organism under discussion are Janus-faced. For example, we can consider an internal state of an ARS_D *either* as an inferentially articulated, conceptually contentful state *or* as an NLO. But (as we have emphasized) it wouldn't be the latter (it wouldn't be a natural *linguistic* object) if it weren't also the former. Nevertheless, we can abstract away from the former role and consider it purely as an item in the natural order, stripped of discursive content.

Nevertheless, many NLOs still have practical import. So do their counterparts in ARS_{AS}. As indicated in our response to (5), various internal states of the animal must be inherently practical (or at least intimately tied to action) in order for those states to picture. However, just as we can abstract away from their discursive character to view NLOs purely as items in the natural order, we can further abstract away from their practical character and view them purely as items standing in isomorphic relations to the extramental items they picture.

Consider an analogy. There is a sense in which a map only passively represents properties and relations and doesn't itself tell you what to do. Thus, a representation of Urbana on a map doesn't by itself tell you to do anything. For Sellars, though, to view the map this way is to view it in abstraction from its function—namely, navigation. Thus, we look at a map (in its primary function) in the context of a plan—say, to travel from Urbana to Chicago, or from Urbana to Indianapolis. And it is such practically freighted cognition that imparts practical significance to the items on the map—and (as we have argued) *in virtue of which we can regard these items as isomorphic with items in the world, and hence as constituting a map in the first place*. Thus, viewing elements on a map as purely descriptive and nonpractical is *parasitic* on viewing them as involved in a web of practical inferences and significances.

Similarly, we can view the internal states of an ARS_A in a similarly abstracted (i.e., purely descriptive) way. There is a sense in which we can say that an internal representation (e.g., of a berry bush) doesn't tell an animal what to do. But to view the representation in this way can only be to abstract from the way that the animal encounters this object in the context of its practical concerns. In the context of the animal being hunted by prey, the representation might have one practical significance ("This way to shelter!"), whereas when the animal is hungry, it might have another ("Here is food!"). But the chief point remains that viewing these states as purely descriptive can only be viewed as an abstraction from these states' practical role, because, as we discussed earlier (section V.B.), it is only because of this internal state's involvement in a complex web of Humean inferences that we are able to talk about the state's role *as* an item on a map—isomorphic with elements in the animal's environment—in the first place. So we can make sense of Sellars's claims in (4) and (5)—as long as we understand that the ability of internal states to play this role is parasitic on their involvement in the animal's practical activity (just as a state's classification as an NLO is parasitic on its Janus-faced role as an inferentially articulated, conceptually contentful state of the agent).

VII. CONCLUDING REMARKS

Although picturing has been an often-underappreciated element, the notion plays a number of absolutely crucial and interrelated roles in Sellars's philosophy. It explains the possibility of semantic content and representation; it explains how mind is integrated in the natural order; it explains (related to the previous two) how there can be an interface between the norm-governed and the purely factual. It demonstrates how representation is tied to practical engagement with the world, and how these two elements—the cognitive and the practical—are inseparable for us

qua rational agents. Thus, the Sellarsian project—which attempts to find a place for norm-governed creatures within the natural order—has picturing absolutely at its core, and picturing cannot be excised from this project without doing fundamental damage to the whole.

Our (perhaps excessively optimistic) goal is to effect a greater unification of Sellars scholarship. As we noted in the introduction, much Sellars scholarship *either* emphasizes the inferentialist, anti-representationalist, and anti-Givenist strand in his thinking *or* emphasizes the role of picturing (and often associates this with Sellars’s *scientia mensura* principle). But to emphasize one and neglect the other is to miss the profound unity of Sellars’s philosophical vision. Philosophers in the first camp must understand that we need picturing to explain how the world rationally constrains our cognitive activity, how there is no mere “frictionless spinning in a void.” But philosophers in the latter camp must recognize that picturing itself is not enough to secure content or representation. Just as important, a narrow focus on picturing elides the crucial difference between ARS_{AS} and ARS_{DS}—both of which picture.

A theory of picturing that neglects the profound differences between Humean ARSs (ARS_{AS}) and Aristotelian ARSs (ARS_{DS}) doesn’t allow for an understanding of how the emergence of rule-governed behavior, or discursive normative practices, introduces a qualitative change in how an animal can picture. The process of learning how to navigate the polydimensional space of reasons transforms how an animal can picture, by liberating us from the shackles of the cognitive constraints due to evolutionary history and giving us the capacity to transform how we picture the world. The key to this, as we see it, is the capacity for counterfactual reasoning and the ability to explicate and evaluate counterfactuals via modal vocabulary.²⁸ Only when *both* picturing *and* the importance of material inference (including

material practical inference) are taken into account can the richness of Sellars's theoretical apparatus be fully appreciated.²⁹

BIBLIOGRAPHY

- Brandom, Robert B. 1979. "Freedom and Constraint by Norms." *American Philosophical Quarterly* 16.3 (July): 187–96.
- Brandom, Robert B. 1994. *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge, MA: Harvard University Press.
- Brandom, Robert B. 2000. *Articulating Reasons: An Introduction to Inferentialism*. Cambridge, MA: Harvard University Press.
- Brandom, Robert B. 2015. *From Empiricism to Expressivism: Brandom Reads Sellars*. Cambridge, MA: Harvard University Press.
- Carroll, Lewis. 1895. "What the Tortoise Said to Achilles." *Mind* 4.14 (April): 278–80. <https://doi.org/10.1093/mind/IV.14.278>
- Dennett, Daniel C. 1995. *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. New York: Simon and Schuster.
- Huebner, Bryce. 2018. "Picturing, Signifying, and Attending." *Belgrade Philosophical Annual* 31: 7–40. <https://doi.org/10.5937/BPA1831007H>
- James, William. (1907) 1979. *Pragmatism: A New Name for Some Old Ways of Thinking*. Cambridge, MA: Harvard University Press.
- Koons, Jeremy Randel. 2019. *The Ethics of Wilfrid Sellars*. New York: Routledge.
- Koons, Jeremy Randel. 2021. "Sellars on Rational Agency as Presupposing Collective Attitudes." In *Groups, Norms and Practices: Essays on Inferentialism and Collective Intentionality*, edited by Ladislav Koreň, Hans Bernhard Schmid, Preston Stovall, and Leo Townsend. Dordrecht, Neth.: Springer, 189–213.
- Kraut, Richard. 2010. "Universals, Metaphysical Explanations, and Pragmatism." *The Journal of Philosophy* 107.11: 590–609. <https://doi.org/10.5840/jphil20101071140>
- Lange, Marc. 2000. "Salience, Supervenience, and Layer Cakes in Sellars's Scientific Realism, McDowell's Moral Realism, and the Philosophy of Mind." *Philosophical Studies* 101: 213–51. <https://doi.org/10.1023/A:1026445113223>
- McDowell, John. 1994. *Mind and World*. Cambridge, MA: Harvard University Press.
- McDowell, John. 1998a. "Intentionality as a Relation." *Journal of Philosophy* 95.9: 471–91. <https://doi.org/10.5840/jphil199895938>
- McDowell, John. 1998b. "Virtue and Reason." In *Mind, Value, and Reality*, edited by John McDowell, 50–73. Cambridge, MA: Harvard University Press.
- Olen, Peter. 2016. *Wilfrid Sellars and the Foundations of Normativity*. London: Palgrave Macmillan.
- O'Shea, James R. 2007. *Wilfrid Sellars: Naturalism with a Normative Turn*. Cambridge: Polity Press.
- Peirce, Charles Sanders. (1868) 1992. "Some Consequences of Four Incapacities." In *The Essential Charles Peirce, Volume I (1867–1893)*, edited by Nathan Houser and Christian Kloesel, 28–55. Bloomington and Indianapolis: Indiana University Press.
- Piaget, Jean. 1952. *The Origins of Intelligence in Children*. New York: International Universities Press.
- Rorty, Richard. 1979. "Transcendental Arguments, Self-Reference, and Pragmatism." In *Transcendental Arguments and Science*, edited by Peter Bieri, Rolf-Peter Horstmann, and Lorenz Krüger. Synthese Library (*Studies in Epistemology, Logic, Methodology, and Philosophy of Science*), vol. 133: 77–103.

- Rorty, Richard. 1988. "Representation, Social Practice, and Truth." *Philosophical Studies* 54.2: 2: 215–28. <https://doi.org/10.1007/BF00354514>
- Rosenberg, Jay. 1974. *Linguistic Representation*. Dordrecht, Neth.: D. Reidel Publishing Company.
- Rosenberg, Jay. 1980. *One World and Our Knowledge of It*. Dordrecht, Neth.: D. Reidel Publishing Company.
- Rosenberg, Jay. 2007. *Wilfrid Sellars: Fusing the Images*. Oxford: Oxford University Press.
- Sachs, Carl B. 2018. "'We Pragmatists Mourn Sellars as a Lost Leader': Sellars's Pragmatist Distinction between Signifying and Picturing." In *Sellars and the History of Modern Philosophy*, edited by Luca Corti and Antonio M. Nunziante. New York: Routledge, 157–177.
- Sachs, Carl B. 2019. "In Defense of Picturing: Sellars's Philosophy of Mind and Cognitive Neuroscience." *Phenomenology and the Cognitive Sciences* 18.4: 669–689. <https://doi.org/10.1007/s11097-018-9598-3>
- Sachs, Carl B. 2022. "A Cybernetic Theory of Persons: How Sellars Naturalized Kant". *Philosophical Inquiries* 10.1: 97–124. <https://doi.org/10.4454/philing.v10i1.389>
- Seibt, Johanna. 1990. *Properties as Processes: A Synoptic Study of W. Sellars's Nominalism*. Atascadero, CA: Ridgeview Publishing Company.
- 2009. "Functions between Reasons and Causes: On Picturing." In *Empiricism, Perceptual Knowledge, Normativity, and Realism*, edited by Willem A. deVries. Oxford: Oxford University Press, 247–283.
- Sellars, Wilfrid. 1963a (SPR). *Science, Perception and Reality*. Atascadero, CA: Ridgeview Publishing Company.
- Sellars, Wilfrid. 1963b (PSIM). "Philosophy and the Scientific Image of Man." In SPR, 1–40.
- Sellars, Wilfrid. 1963c (BBK). "Being and Being Known." In SPR, 41–59.
- Sellars, Wilfrid. 1963d (EPM). "Empiricism and the Philosophy of Mind." In SPR, 127–196.
- Sellars, Wilfrid. 1963e (TC). "Truth and 'Correspondence.'" In SPR, 197–224.
- Sellars, Wilfrid. 1963f (SRLG). "Some Reflections on Language Games." In SPR, 321–58.
- Sellars, Wilfrid. 1967 (SM). *Science and Metaphysics: Variations on Kantian Themes*. The John Locke Lectures for 1965–66. Atascadero, CA: Ridgeview Publishing Company.
- Sellars, Wilfrid. 1973 (AAE). "Actions and Events." *Noûs* 7.2 (May): 179–202. <https://doi.org/10.2307/2214491>
- Sellars, Wilfrid. 1980a (NAO). *Naturalism and Ontology*. Atascadero, CA: Ridgeview Publishing Company.
- Sellars, Wilfrid. 1980b (PPPW). *Pure Pragmatics and Possible Worlds: The Early Essays of Wilfrid Sellars*, edited by Jeffrey F. Sicha. Atascadero, CA: Ridgeview Publishing Company.
- Sellars, Wilfrid. 1980c (LRB). "Language, Rules, and Behavior." In PPPW, 129–155.
- Sellars, Wilfrid. 1981 (MEV). "Mental Events." *Philosophical Studies* 39: 325–45. <https://doi.org/10.1007/BF00360342>
- Sellars, Wilfrid. 2017 (WSNDL). *Wilfrid Sellars: Notre Dame Lectures, 1969-86*, edited by Pedro Amaral. Atascadero, CA: Ridgeview Publishing Company.
- Shapiro, Lionel. 2011. "Intentional Relations and the Sideways-On View: On McDowell's Critique of Sellars." *European Journal of Philosophy* 21.2: 300–319. <https://doi.org/10.1111/j.1468-0378.2010.00448.x>

- Stovall, Preston. 2016. "Nature, Purpose, and Norm: A Program in American Philosophy." *Journal of the American Philosophical Association* 2.4: 617–36. <https://doi.org/10.1017/apa.2016.32>
- Stovall, Preston. 2022. *The Single-Minded Animal: Shared Intentionality, Normativity, and the Foundations of Discursive Cognition*. New York: Routledge.
- Tolman, Edward C. 1948. "Cognitive Maps in Rats and Men." *Psychological Review* 55.4: 189–208. <https://doi.org/10.1037/h0061626>
- Tomasello, Michael. 2014. *A Natural History of Human Thinking*. Cambridge, MA: Harvard University Press.

¹ For more on this strand of naturalism in Sellars's thinking, see Stovall (2016).

² This is, of course, precisely what McDowell objects to in his third Woodbridge lecture, "Intentionality as a Relation" (McDowell 1998a). But for a contrary interpretation of Sellars, see Shapiro (2011).

³ For a detailed study of Sellars's nominalism, see Seibt (1990). Sellars doesn't deny that *in a sense* there are abstract entities, such as *triangularity*. But these aren't part of the ultimate "furniture of the universe"; see Kraut (2010).

⁴ Meaning does, of course, require relations *between embodied minds*. This is why Sellars's practical philosophy becomes crucial to understanding his semantics.

⁵ Sellars's phrase is from TC, p. 212.

⁶ Cf. "the pseudo-psychology of cognitive givenness," LRB p. p. 305/§26.

⁷ For the importance of cybernetics for Sellars's philosophy of mind, see Sachs (2022).

⁸ Although Rosenberg accepted picturing early on (1974), he soon rejected it (1980). Rorty's criticism of Sellarsian picturing (1979) might have played a role in Rosenberg's change of mind. For a criticism of Rorty's argument, see Sachs (2018). Much later, Rosenberg (2007) accepts Sellarsian picturing; for further developments of Sellarsian picturing based on Rosenberg, see Seibt (2009), Huebner (2018), and Sachs (2019). See also Brandom's account of the (mostly unproductive) conversations he had with Sellars concerning picturing: "I was never able to understand how he understood such matter-of-factual picturing relations to be related to the normatively characterizable discursive practices that (he and I agreed) alone deserved to be thought of in genuinely *semantic* terms. Eventually he confessed that he did not take himself clearly to understand the relation, either" (Brandom 2015, p. 13).

⁹ Or at least "intensional" states, in the case of non-concept-using animals, which represent objects "as" having properties in "proto-conceptual" ways, i.e., not *strictly* intentional or conceptual ways, as Sellars reserves those terms for logical representers. Jim O'Shea suggested this qualification.

¹⁰ This is a generalization, and as such it does not picture. Generalizations are formulated in the metalanguage, and we will see later in this essay the crucial role that the metalanguage plays in the constitution of conceptual representational systems such as ourselves. For now, we will merely note that the examples that follow involve temporal sequences that picture (since they do not involve generalizations).

¹¹ However, Sellars also wants to extend his account of picturing to theoretical picturing, as he does in SM V. Theoretical picturing also requires the relevant language-entry moves (e.g., from a certain perceptual input to the judgment, "Lo, an electron!").

¹² It should be noted that Sellars's criticism of the elision between "the normative" and "the natural" turns on claims about being able to analyze normative facts into natural facts. Given the Russell/Moore conception of conceptual analysis that Sellars accepted, Sellars's insistence on the irreducibility of the normative to the natural means that there are no analytic sentences, "true by meaning alone," that have normative facts on the right-hand side and natural facts on the left-hand side.

¹³ Bill deVries has asked about statements in which a predicate occurs within the context of negation, such as "The stop sign is not green." His suggestion is that such statements do not picture because only atomic statements picture. This may be Sellars's view (although Rosenberg [1974, chapter 7] has an account of how even complex statements can picture). But even if we hew to deVries's suggestion, negated sentences ("The stop sign is not green," "No stop sign is green," etc.) play an important role in picturing in that (even if they themselves do not picture) they inhibit the development or maintenance of faulty isomorphisms.

¹⁴ A good deal depends on “external to what, exactly?” Rorty (1979) objects to picturing on the grounds that it would be external to all human conceptual activity, *sub specie aeternitatis*, a shadow of God to be expunged; see also Rorty (1988). For a contrary view that picturing is part of the scientific image and thus external to the manifest image specifically and therefore *not* external to all human cognition, see Sachs (2018).

¹⁵ The arguments of this section are presented in much greater detail in chapters 6 and 14 of Koons (2019).

¹⁶ Sellars explicitly cites the argument made “most brilliantly by Lewis Carroll” in AAE (p. 201, note 2).

¹⁷ This understanding of cognition as saturated by practical concerns is indebted to McDowell, especially McDowell (1998b); for an analysis of the essentially Sellarsian reasoning at work in McDowell’s position, see Lange (2000). For a reconstruction of Lange’s argument and a more general application of it to the relation between the scientific and manifest images, see chapter 2 of Koons (2019).

¹⁸ We may interpret a differential responder as representing an object as being *somehow*—say, a thermometer as representing the temperature as 72°F. But its representation *as* is wholly parasitic on our capacity to represent *as*; it merely produces an output that (within a narrow range) corresponds with a particular input. That a differential responder’s representational ability is wholly parasitic on our own is made clear by the fact that a trained parrot does not represent red objects *as* red; nor does a piece of iron represent the environment *as* containing oxygen. That we can put differential responders to work within a representational system merely reflects the fact that we are representers.

¹⁹ For an in-depth development of this argument and an application of it to Sellars’s metaethical framework, see Koons (2019), chapters 10 and 11.

²⁰ This is Kantian in the following sense: We apply the concepts of the understanding to the products of receptivity, and so any reordering/restructuring of the concepts of the understanding can alter the cognitive character of our intuitions (even if nothing changes about causal inputs received from sensibility).

²¹ The concept of representation-entry transitions and representation-exit transitions is based on what Sellars calls language-entry transitions and language-exit transitions (see SRLG) but generalized for nondiscursive animal representational systems.

²² We are of course eliding a number of complications here. On the “default and challenge” model common to pragmatists (compare, for example, James’s “credit system” [James 1907/1979, Lecture VI, p. 100] or Peirce on “real doubt” [Peirce 1868/1992, pp. 28-9]), not just any challenge requires a response, as a person with a *prima facie* entitlement to a claim may often ignore “bare challenges.” But the role that the metalanguage plays in this is essential: It allows for such challenges, thus allowing for revision to the conceptual framework that is both (a) *rational* [i.e., reason-guided] and (b) *interpersonal*.

²³ However, we do not wish to commit ourselves to any theory of animal learning.

²⁴ Another way of putting this point is to say that what birds consume is not liable to feedback correction from other birds.

²⁵ With respect to our view that practical picturing is crucial for social coordination, we learned (shortly after this paper was accepted for publication) that Preston Stovall (2022) develops very similar views; he calls his particular conception “shared practical picturing.”

²⁶ Tying together a number of themes here, the ability to use a metalanguage to represent inferential proprieties in the intersubjective game of giving and asking for reasons is key to the evolutionary function of logic: logic evolved to facilitate collective action by getting everyone to agree on what to believe and therefore on what to do.

²⁷ The reference is to Daniel Dennett (1995).

²⁸ It is beyond the scope of the current paper to show how the account of picturing developed here also helps make sense of the role of picturing in the self-correcting project of empirical science.

²⁹ We are grateful to the participants in the International Sellars Colloquium (organized by Luz Christopher Seiberth) for offering us helpful feedback on a draft of this essay. An anonymous referee for the *Journal of Philosophical Research* also provided us with detailed feedback that helped us clarify our central argument.