Sellars's Two Responses to Skepticism

Griffin Klemick
Hope College

This is a post-peer-review, pre-copyedit version of an article forthcoming in *Synthese*.

ABSTRACT: This paper offers a critical interpretation and evaluation of Wilfrid Sellars's treatment of skepticism about empirical justification. It defends three central claims. First, against the suggestion that Sellars's work simply bypasses traditional skeptical problems, I make the novel interpretive claim that Sellars not only addresses skepticism about empirical justification, but offers two independent (albeit sketchy) arguments against it: a transcendental argument that the likely truth of our perceptual beliefs is a necessary condition of the possibility of empirical content, and a pragmatic argument that we're warranted in accepting their likely truth in virtue of our aim of being effective agents. To the extent these have previously been distinctly formulated by commentators, the transcendental argument has been regarded as forceful, while the pragmatic argument has been dismissed as non-responsive. My second and third claims challenge this understanding. I argue, second, that examination of the literature relating to transcendental arguments from semantic externalism like Sellars's (especially concerning the McKinsey paradox) suggests that such arguments are unpromising, while, third, a modified version of his pragmatic argument represents a powerful skeptical solution to skepticism about empirical justification, one that answers the worry that such skepticism would undermine the rationality of all our practical commitments.

1. Introduction: Who Cares what Sellars Thought about Skepticism?

Skepticism about empirical justification is (naturally enough) the thesis that none of our empirical beliefs is justified. In this paper I offer a critical interpretation of Wilfrid Sellars's response to this thesis—or, better, of Sellars's *responses*, since my central interpretive claim is that Sellars actually offers two independent arguments against skepticism, a transcendental argument and a purely pragmatic argument. These haven't previously been clearly

distinguished by commentators, but to the extent they've been identified in Sellars's writings, the present consensus seems to be that his transcendental argument is forceful and promising, while his pragmatic argument is weak or even downright non-responsive to the skeptical problem. My central substantive claim is that this consensus is doubly mistaken: prospects for transcendental arguments of the general shape of Sellars's are actually rather dim, while his pragmatic argument, while admittedly (but, in my view, appropriately) modest in its ambitions, represents a promising response to skepticism about empirical justification.

Why does Sellars's position regarding skepticism merit attention? Well, obviously, if his position is sound, we'll learn something about skepticism by doing so. But further, in terms of the payoff for our understanding of Sellars's philosophy, I see two key reasons. The first is that, while Sellars's epistemology has recently garnered little detailed attention in comparison to, say, his philosophies of language and mind,[1] interest in these latter topics has often been motivated by the promise of an epistemological payoff. For instance, Richard Rorty, who probably did more than anyone to rescue Sellars's work from obscurity, was interested in Sellars chiefly because he took Sellars's argument that the Given is a myth to dissolve classical epistemological problems. In Rorty's view, Sellars's argument sweeps away "*tertia*" traditionally posited to mediate between humans and the world, thereby freeing us from a representationalism about perception and content that renders skepticism unanswerable.[2] And following Rorty, other prominent philosophers with Sellarsian influences and affinities have embraced forms of epistemological quietism or doxastic

[1] This is natural: Sellars wrote many fewer essays directly concerning epistemology than concerning language and mind, and the former often rely on more fundamental theses concerning the latter topics.

[2] On the roots of Rorty's opposition to representationalism in his efforts to prevent our becoming "patsies" for skepticism, and his esteem for Sellars's work in that light, see Brandom 2013: 92–93. Regarding his consequent rejection of "*tertia*"—albeit with reference to Davidson rather than Sellars—see Rorty 1991: 138–39. Compare Levine's (2019: 1) claim that, for Rorty, rejecting the Myth of the Given with Sellars enables us to "dissolv[e] the epistemological problematic" and cease trying to "bridge the divide between mind and world."

conservatism, dismissing rather than answering the skeptic's challenge to independently defend our entitlement to hold our empirical beliefs.[3] On my interpretation, however, Sellars did not aim simply to dissolve skeptical problems: he saw that he needed a robust answer to broadly Humean arguments for skepticism about empirical justification, and he sought to provide one. My interpretation thus has the potential to challenge not only alternative readings of Sellars's epistemology, but more fundamentally the prevalent impression that his philosophy bypasses epistemological problems altogether and is important in significant part because it enables us to do likewise.

The second reason for considering Sellars's position regarding skepticism is that it's essential for the assessment of his scientific realism. As congenial as pragmatist opponents of realism like Rorty found his philosophies of perception and content, "Sellars's deepest philosophical commitment is to naturalism" (deVries 2005: 15), and he saw it as essential that his broadly pragmatist, non-representationalist theory of content should be rendered consistent with realism about the natural world and about the capacity of the natural sciences finally to yield knowledge of it as it really is. And lest this theory of content land him in a linguistic idealism that precludes this realist stance, he introduced his theory of picturing. (Picturing is a sort of complex co-variation between our assertions and worldly items that constitutes a non-semantic isomorphism between nexuses of the former—conceived purely naturalistically, i.e. in abstraction from the norms of use that, for Sellars, constitute their

[3] Regarding quietism, see McDowell 1994/1996: xiii–xiv, 143; 2009: 206. (Davidson's [2001: 154] famous concession to Rorty that we shouldn't answer the skeptic but only tell him to "get lost" may seem quietist, but Davidson actually aims to entitle himself to dismiss the skeptic thus via a transcendental argument from semantic externalism much like Sellars's: "It should now be clear what insures that our view of the world is, in its plainest features, largely correct. The reason is that the stimuli that cause our most basic verbal responses also determine what those verbal responses mean, and the content of the beliefs that accompany them" [ibid.: 213].) Regarding conservatism, see the "default-and-challenge" model of justification proposed by Brandom (1994: 176–79) and endorsed by Williams (2009: §7).

meanings—and nexuses of the latter.[4]) The point of this theory was to secure the place of our thought and talk in the natural world and, given his definition of the truth of basic empirical assertions in terms of their correctly picturing (*SM* V: ¶¶9, 57),[5] the accountability of our empirical discourse to how things stand in it. Rorty objects, however, that the appeal to picturing represents a backsliding into representationalism that renders skepticism about empirical justification unanswerable: since our assertions' success at picturing isn't transparent to us, our theories might leave us "predicting better and better while picturing worse and worse" (1991: 155), which possibility deprives us of entitlement to regard their enabling predictive success as a mark of their truth or as warranting us in believing them. Now, this objection is hardly dispositive. It's not specific to Sellars, after all: it applies equally to any account of empirical meaning on which empirical truth isn't epistemically constrained. And since proponents of such accounts have offered a myriad of responses to skepticism, Rorty's objection will succeed only if every such response fails. Still, the objection highlights skepticism's importance for the assessment of Sellars's core philosophical commitments: his defense of his scientific realism was fully adequate only if his response to skepticism was adequate, too. Anyone sympathetic to the former thus should also be interested in the latter.

In §2 I trace Sellars's engagement with skepticism back to "Empiricism and the Philosophy of Mind," arguing that Sellars fails in his attempt there to respond to Humean skeptical worries without succumbing to coherentism. In later writings, though, he breaks with coherentism by enlisting certain epistemic principles—centrally, for our purposes, the principle that our perceptual judgments are likely to be true—as foundations of empirical

[4] Given the theory of picturing's complexity, we cannot survey it in detail here. (See deVries 2019 for a "primer.") What's important for our purposes is simply that the theory yields a realist theory of empirical truth that raises skeptical worries.

[5] I refer to Sellars's works by standard abbreviations: see the Bibliography, §1.

justification. His two responses to skepticism are two independent arguments in defense of our accepting this principle: one is a transcendental argument, rooted in semantic externalism, that the likely truth of our perceptual beliefs is a necessary condition of the possibility of empirical content, and the other is a pragmatic argument that we're warranted in accepting that they're likely to be true, given our aim of being effective agents. The remaining sections critically assess these responses. Since Sellars himself states them only briefly and never seriously undertakes to defend them, this requires that we go beyond his own writings somewhat, especially in §3, where I enlist the literature on semantic externalism's anti-skeptical import and its relationship to privileged access to mental contents—principally, the literature surrounding the McKinsey paradox—in an argument that transcendental arguments like Sellars's are unlikely to succeed. In §4 I turn to the pragmatic argument, briefly motivating modifications that, in my view, enable the argument to succeed on its own terms, and arguing that its success represents a significant achievement: while it isn't a direct response to skepticism about empirical justification, it is a persuasive argument that such skepticism does not undermine our warrant to reason and act in conformity with our perceptual evidence—and therefore doesn't undermine the reasonability of our practical commitments—as one might naturally have feared it would.

2. Sellars's Path to the Two Responses

The variety of skeptical problem that Sellars most frequently engages is not Cartesian skeptical scenarios, but rather Humean worries about rule-circularity.[6] This makes sense,

[6] In the first *Enquiry*, §IV, Hume famously argues that empirical justifications of the principle that induction is reliable covertly affirm the principle *ab initio*. Actually, Sellars had already offered a powerful answer to that particular skeptical worry two years prior to delivering the lectures that would become EPM, pointing out that his inferentialist theory of meaning yields the consequence that "the notion of a language which enables one to state empirical matters of fact, but contains no empirical moves is . . . chimerical," and thus that the "classical

given the way he handles a problem that arises from his rejection of the Given as a myth. The "most straightforward form" of the Myth, he writes, is to posit "nonverbal episodes of awareness . . . *that* something is the case, e.g., *that this is green*—which . . . have an intrinsic authority (they are, so to speak 'self-authenticating')" and which our perceptual judgments and observation reports express (EPM: §VIII/¶34).[7] Having denied that there are any such self-authenticating experiential states to imbue such judgments and reports with epistemic authority, Sellars immediately faces the question what alternative source their authority could have.

Part of his answer is reliabilist, and part appeals to social practice. One part of the reason our tokens of (e.g.) 'This is green' express perceptual knowledge (when they do) is that they manifest a disposition to produce such tokens (aloud or merely in thought) just when the speaker is looking at something green in standard conditions. And another part is that they result from training by one's linguistic community and constitute "instance[s] of a general mode of behavior which, in a given linguistic community, it is reasonable to sanction and support" (EPM: §VIII/¶35). This suffices to ground the evaluation of observation reports as correct or incorrect. But to treat an observation report not merely as correctly made but as expressing *knowledge* requires, further, that the report's authority "must *in some sense* be recognized by the person whose report it is" (ibid.):

> for [an observation report] "This is green" to "express observational knowledge," not only must it be a *symptom* or *sign* of the presence of a green object in standard

---

'fiction' of an inductive leap which takes its point of departure from an observation base undefiled by any notion as to how things hang together is not a fiction but an absurdity" (SRLG: ¶63). But this response—that to be justified in asserting any empirical statement is to be justified in inferring other matters of fact on its basis—does not solve, but raises the broader worry about rule-circularity with which this section is concerned: namely, how can we be justified in asserting particular empirical statements without prior warrant to assert that the faculties by which we learn them are reliable, and vice-versa?

[7] Sellars frequently numbers the sections of his papers using Roman numerals and their paragraphs using Arabic numerals. When citing his works, I always include the latter; due to EPM's prominence, I include the former, too, when citing it.

> conditions, but the perceiver must know that tokens of "This is green" *are* symptoms of the presence of green objects in conditions which are standard for visual perception. (ibid.).

This "epistemic reflexivity requirement," as Willem deVries (2005: 277) calls it, has the effect of thoroughly excising the Myth of the Given from Sellars's account of perceptual knowledge. On this account, observational knowledge isn't self-authenticating: it doesn't "stand on its own feet," as it would have to in order to provide a foundation for all empirical knowledge (EPM: §VIII/¶36). Rather, instances of observational knowledge of particular matters of fact presuppose knowledge of general facts about the specific conditions in which particular perceptions reliably indicate particular perceivable states of affairs.[8] But this requirement raises the specter of Humean skepticism, too. After all, how could we possibly come to know these general facts except on the basis of particular perceptual observations? If there's no other way to do so, then perceptual knowledge will be impossible to acquire: we would have to have particular items of perceptual knowledge prior to general knowledge about when perception is reliable, and vice-versa, trapping us in a vicious circle.

In EPM, §VIII, Sellars offers a response to this problem of circularity that I take to be coherentist. On this response, while it's true that a subject must *now* know in what conditions perception is reliable to count as *now* knowing via perception that particular objects are green, and while having these items of general knowledge now requires the prior acquisition over time of linguistic habits to make (proto-conceptual counterparts of) observation reports, her having these items of general knowledge now doesn't require linguistic episodes in that acquisition process to have antecedently counted as knowledge

[8] I reject Brandom's (1997: 153) claim that, in EPM, §VIII, Sellars objects only to classical empiricism's "hierarchical picture of *understanding*," not its "hierarchical picture of *justification*." On Sellars's position there, our observation reports depend on our general empirical knowledge not only for their meaningfulness, but also for their justification. Cf. deVries and Triplett 2000: 98–99, 107. (In later works, we'll see, Sellars does endorse a hierarchical picture of empirical justification—by dropping the idea that the epistemic principles fundamental to empirical knowledge epistemically depend on particular perceptual judgments in any respect.)

(prior to receiving support in turn from the items of general knowledge). Rather, Sellars would say with the later Wittgenstein, "light dawns slowly over the whole": neither the general nor the particular beliefs count as knowledge prior to support by the other, but the status of the former beliefs as knowledge emerges jointly with that of the latter as they come to support one another epistemically (¶37).

Sellars's dissatisfaction with coherentism was nevertheless already clear in EPM from his remarks that "to say that empirical knowledge has *no* foundation . . . is to suggest that it is really 'empirical knowledge so-called,' and to put it in a box with rumors and hoaxes" (EPM: §VIII/¶38), and that it won't do to embrace "the picture of a great Hegelian serpent of knowledge with its tail in its mouth (Where does it begin?)" (ibid.). But I don't see that EPM, §VIII succeeds in distancing his account of empirical justification substantially from coherentism:[9] the epistemic reflexivity requirement seems to prevent observation reports' being autonomously justified so as to serve as even a defeasible foundation for empirical warrant. (They may be immediately *caused*, independently of one's other beliefs—but they aren't at all immediately *justified*.) While he does suggest that there must be "*some* point" to the picture of empirical knowledge as resting on observation reports that aren't reciprocally

[9] Noting Sellars's concluding remarks that "*Above all*, the [traditional epistemological] picture is misleading because of its static character," and that not only it but also Hegelian coherentism is unsatisfactory given that "empirical knowledge . . . is rational . . . because it is a self-correcting enterprise which can put *any* claim in jeopardy, though not all at once" (EPM: §VIII/¶38), Williams suggests that Sellars departs from coherentism by "insist[ing] that warrant be viewed dynamically, hence historically: our world-picture is warranted because it is the result of the rational correction of a previous picture" (2009: 175). But it's not clear why this should render our present world-picture warranted, absent any antecedent warrant for the previous picture. (If I render an incoherent fairy tale coherent, this is a rational correction, but it doesn't justify me in believing the resulting contents!) And if this antecedent warrant is grounded simply in the previous picture's coherence, then the account's difference from coherentism is superficial. (Certainly, this seems true, anyway, if the measure of the rationality of a given putative correction is simply that it enables "increasing explanatory coherence," as Williams suggests [ibid.: 183].)

dependent on other justified empirical beliefs (ibid.), then, it isn't clear that his EPM account actually made good on this suggestion.[10]

We thus shouldn't be surprised either that Sellars continued to grapple with epistemic circularity in later writings or that he did so by revising his EPM position to sharply break with coherentism. Given the suggestion just considered, though, we may well be surprised regarding which particular beliefs he enlisted to serve as foundations of empirical warrant. For on his later view, while the justification of observational beliefs rests in some sense on the soundness of an inference from certain special epistemic principles, these principles (despite being empirical claims themselves) don't likewise depend on particular observational beliefs for their justification.

Of course, an observational belief is "a *paradigm* case of non-inferential belief" in one sense: it's non-inferentially *elicited*. Indeed, Sellars further suggests: "*The authority of [such] thinking accrues to it in quite a different way*" from the way authority accrues to conclusions validly inferred from true premises (SK II: ¶37). Still, he notes, given our knowledge that mature human beings' spontaneous observation reports are generally reliable guides to perceptible worldly states of affairs, if we heard someone making such a report and knew that normal perceptual conditions obtained, we'd be in a position to infer that there's good reason to believe that things are as the report describes them. Indeed, I can perform such inferences with respect to my own observational beliefs, concluding that they're justified by appeal to the character and context of my perception and the general reliability of perceivers

[10] DeVries and Triplett rightly note that the epistemology Sellars sketches in EPM, §VIII isn't "a traditional foundationalism," but likewise claim that it isn't "a traditional coherentism," since observation reports "*do* have a different epistemic status . . . from other empirical beliefs" (2000: 107). I'm unsure how this proves it isn't coherentist, though, given that this different status, whatever exactly it is, doesn't involve their being epistemically independent. (On the preceding page, deVries and Triplett are more circumspect, suggesting that Sellars needed to say more about "the special role of observation reports" to make it clear where his account fell on the foundationalism–coherentism spectrum. I prefer to say that, having failed to supply an adequate account of this "special role," Sellars hadn't succeeded yet in breaking with coherentism.)

in such conditions. My observational beliefs needn't be *based* on conscious inferences of this sort to count as justified, nor need I subsequently make such an inference with respect to every observational belief of mine for that belief to count as justified. Still, Sellars maintains, were I not in a position to make such inferences concerning any of my observational beliefs, I would lack the knowledge of my perceptual reliability necessary for such beliefs to count as justified. In this sense, at least, he's prepared to assert that "the justification of [a perceptual belief] is an inferential justification" (SK III: ¶35)—and so, "in effect, that all justification is inferential" (*WSNDL*: 157).[11] For even my justification for non-inferentially elicited perceptual beliefs rests on my justification to accept the general claim that such beliefs, formed in normal perceptual conditions, are likely to be true, and my ability to reason from this premise to the probable truth of the perceptual belief in question.

The obvious question, then, is how I can be justified in believing this general claim.[12] Even if one ground for accepting that our perceptual beliefs are likely to be true is that this claim seems to be empirically well-supported, we need an independent ground for accepting

---

[11] Sellars's position thus is naturally summarized as the thesis that our perceptual beliefs are only non-inferentially elicited, not also non-inferentially justified (cf. deVries 2005: 296n51). But I think his claim that our justification for perceptual beliefs is inferential risks being misleading, since it's naturally interpreted as entailing that such justification must be *based on* our actual performance of an inference, and he doesn't in fact maintain that this is necessary. I think he would have framed his position more straightforwardly had he simply said that, on his view, the justification my perceptual beliefs have wouldn't be possible but for my knowledge of the availability of an inference schema from my perceptual reliability in standard conditions to my justification to hold particular perceptual beliefs. My warrant to believe in such reliability is thus epistemically prior to my warrant for any particular perceptual belief, even if the latter doesn't always derive from the former by explicit inference.

[12] Readers familiar with Sellars might wonder (notwithstanding the clear evidence of MGEC: §IV) whether he really aims to warrant us in accepting that our perceptual judgments are likely to be true—for doesn't he deny that they *are* likely to be true? After all, Sellars holds that the commonsense, Manifest Image conceptual framework in which we make our ordinary perceptual judgments is "radically false, i.e., there *really* are no such things as the physical objects and processes" it posits (SRI: ¶48). Aren't all such judgments false, then? Not for Sellars, for two reasons. First, since "truth in the 'absolute' sense is . . . relative to *our* language" (*SM* V: ¶48), when we Manifest Image-users take our ordinary perceptual judgments as true, we're typically correct to do so. Second, true propositions in our own conceptual framework have counterpart propositions that perform a sufficiently similar functional role in an ideally adequate scientific conceptual framework, and those counterparts are true with respect to the latter framework (ibid.: ¶¶73–74).

this to avoid falling prey to the Humean charge of circularity (MGEC: ¶66). What might this independent ground be?

Though Sellars himself doesn't explicitly say this, I claim that he actually offers two independent responses to this question: one could accept either argument while rejecting the other. The responses differ both in their starting points and in their modes of argument. The first is a *transcendental* argument: it begins from the premise that our thoughts and assertions are empirically contentful, and it argues that a necessary condition of the possibility of the truth of this premise is that our perceptual beliefs are likely to be true. The second is a *pragmatic* argument: it begins from the premise that we have the goal of being effective agents, and it aims to establish that it's practically reasonable for us to accept that our perceptual judgments are likely to be true on the grounds that the truth of this claim is necessary for our realization of this goal. I'll offer an uncritical statement of each argument in this section before fleshing out and critically examining each in subsequent sections, arguing that the former's prospects of success are poor while a promising version of the latter is available.[13]

[13] I'm not aware of other commentators who've seen two independent arguments at work in Sellars's reflections on this topic. DeVries (2005: 139–40) stresses Sellars's pragmatic argument in expounding MGEC's conclusion; he quotes the passage in which Sellars offers the transcendental argument, but he doesn't explicitly identify it as such or distinguish it from the pragmatic one. O'Shea, by contrast, in one treatment (2007: 131–35) discusses the transcendental argument in depth but assimilates the pragmatic dimension of Sellars's view to it, and in a later one (2011: §§IV–V) draws on elements of both, but without noting their independence. Williams (2009: 173–74) perhaps comes closest to distinguishing them, considering a reading on which Sellars offers a "merely strategic or pragmatic argument" and then replying that, actually, he offers a stronger argument "akin to a transcendental argument." But given that the elements of his view that lend themselves to each reading are independent of one another, we should say instead that he offers both arguments, though without carefully distinguishing them—and, indeed, perhaps regarding them as more closely connected than they really are.

First, Sellars remarks that, if asked why it's reasonable to accept the claim that our perceptual[14] beliefs are likely to be true, he would argue that these beliefs

> are elements in a conceptual framework that defines what it is to be a finite knower in a world one never made. . . .
>
> To be one who makes epistemic appraisals is to be in this framework. And to be in this framework is to appreciate the interplay of the reasonableness of inductive hypotheses and of [perceptual] judgments. (MGEC: ¶¶73, 75).

As he summarizes this line of thought elsewhere: "*We have to be in this framework to be thinking and perceiving beings at all*" (SK III: ¶45). Since we manifestly are thinking and perceiving beings, we must be justified in believing that our perceptual judgments are likely to be true.

On this view, to epistemically appraise our empirical beliefs is already somehow to be committed to the likelihood of our perceptual beliefs to be true. Since one must engage in this activity to argue for skepticism about empirical justification, this argument, if it could be made good, would establish that even the skeptic is committed to the claim that our perceptual beliefs are justified. Moreover, since, on Sellars's view, grasping the meaning of an empirical claim is a matter of grasping the conditions under which it's correctly assertible (together with the conditions under which it renders further statements correctly assertible), one must engage in this activity even to understand empirical statements in the first place. As James O'Shea (2007: 132) argues, then, we can see this strand of Sellars's response to the skeptical challenge as a transcendental argument for the likely truth of our perceptual beliefs, one based on so slender a premise as the empirical contentfulness of our thought and language—or, as Sellars himself liked to put the point, based simply on our being "organism[s] whose language is *about* the world in which it is *used*" (SK III: ¶46). Sellars's first anti-skeptical argument thus amounts to this:

---

[14] Sellars himself discusses our *IPM*—introspective, perceptual, and memory—*judgments*. For simplicity and continuity with the preceding discussion, I'll apply the arguments only to perceptual judgments or beliefs, though Sellars's application of them to introspective judgments, too, will be relevant to my argument in §3.

The Transcendental Argument

A. My thoughts and assertions are empirically contentful.
B. My thoughts and assertions can be empirically contentful only if my perceptual judgments are likely to be true.
C. Thus, my perceptual judgments are likely to be true.

This doesn't get us very far, however, if we cannot find some *prima facie* motivation for (B). Few would venture to deny that their thought and talk are empirically contentful. (Indeed, at least at first blush, this seems self-undermining.) But why think that this presupposes the likely truth of their perceptual judgments?

To see how Sellars answers this question, let's return to his account of the sort of authority that accrues to perceptual beliefs, a sort different from the kind that attaches to beliefs arrived at via sound arguments. This authority, Sellars remarks, "can be traced to the fact that [the speaker] has learned how to use the relevant words in perceptual contexts" (SK II: ¶37; italics removed). What does this mean?

On Sellars's theory of meaning as normative functional role, an empirical statement's meaning is given by the rules governing its material-inferential relations to other statements, as well as those describing the perceptual situations in which its assertion is warranted and the practical activity its assertion warrants. Now, some such statements—those concerning unobservable entities—won't relate to perceptual and practical situations directly, but only via the mediation of other statements.[15] But the meaning of a statement with an observational use—for example, 'this is a dog'—is partly constituted, not only by the statement's being properly inferred from 'this is a goldendoodle' and in turn licensing the inference of 'this is an animal,' but further from its being properly asserted when the speaker encounters a dog—and not when she encounters a fox or raccoon—in standard perceptual

[15] That said, for Sellars, the unobservable/observable distinction is methodological, not ontological, and scientific progress enables us to deploy observationally terms that previously had only a theoretical use.

conditions. And since, as Sellars asserts in what O'Shea (2007: 50) has dubbed his *norm/nature meta-principle*, espousal of principles is reflected in uniformities of performance (TC: 216), a person's commitment to follow these rules entails that, *ceteris paribus*, she actually does so. This can also be seen from the fact that, as Sellars remarks here, she has been taught the meanings of perceptual terms through training in their use in perceptual contexts, instilling in her the dispositions to use them in just the correct circumstances. Accordingly, on Sellars's account, that our body of observation reports results from a set of reliable dispositions is necessary not only for it to express perceptual knowledge, but for it even to count as empirically meaningful. This is the core of Sellars's transcendental argument for the likelihood that our perceptual judgments are true. We'll assess its prospects in §3.[16]

In one text in which Sellars offers this argument, though, he doesn't stop there. He further appeals to "the necessary connection between being in the framework of epistemic evaluation and being agents," arguing that this connection is the ultimate justification for our accepting that our perceptual judgments are likely to be true. For if such judgments aren't likely to be true, "the concept of effective agency has no application," since "agency, to be effective, involves having reliable cognitive maps of ourselves and our environment," and for our perceptual judgments to be (likely to be) "correct mappings of ourselves and our

[16] Why think this argument presupposes semantic externalism, as I have suggested repeatedly above? Because "externalism is the claim that, at least in some cases, . . . the possession of a mental state with a given content is environmentally dependent," in that it wouldn't be "possible for an individual to entertain this content unless the environment contain[ed] the requisite objects" and other entities the content denotes (Rowlands et al. 2020: §3). And the central thrust of Sellars's argument is that, given the norm/nature metaprinciple, I couldn't grasp the content of a perceptual judgment like 'This is a dog' if I didn't reliably succeed in making it, *ceteris paribus*, just when perceptually confronted by a dog. But then I couldn't grasp this content unless my environment contained dogs. (Again, this isn't to say that Sellars is committed to the existence of dogs *as conceived within our commonsense/Manifest Image conceptual scheme*: "That there is no such thing as O conceived in the framework of common sense, is compatible with the idea that there is such a thing as O conceived in another framework, [namely,] that of physical theory" [*SM* V: ¶64; recall my footnote 12]. Compare Matsui's discussion of Sellars's *ideal successor externalism* and the preservation of reference across conceptual change that it involves [2021: esp. 141–42].)

circumstances" just is for them to be (likely to be) true (MGEC: ¶¶80–83; cf. *SM* V: ¶9). Our being effective agents thus presupposes that our perceptual judgments are likely to be true.[17]

This thought, as I noted, occurs intermingled with Sellars's transcendental argument. Even so, it constitutes an independent anti-skeptical argument. It isn't really a transcendental argument, since an effective transcendental argument must address skepticism about empirical justification, and so must proceed from premises we can know independently of empirical investigation. The claim that our thought and language are empirically contentful would appear to be a strong candidate for meeting this bar. But the claim that we're effective agents isn't obviously a strong candidate for this status: the question whether our actions will tend reliably to achieve our desired ends seems obviously an empirical matter in a way that the question whether our thought and language are so much as about the world doesn't. So, the appeal to agency doesn't plausibly found a transcendental argument for the likely truth of our perceptual beliefs.

Instead, it seems to found a purely *pragmatic* argument for our warrant to regard them as likely to be true. And this is suggested by the way Sellars initially leads into his response to the skeptic in this context:

> I think that [a way in which it could be *independently* reasonable to accept that our perceptual judgments are likely to be true][18] can be found by following . . . [an]

---

[17] While Sellars's initial allusion to "the necessary connection between being in the framework of epistemic evaluation and being agents," then, might naturally have led one to expect an extension of his transcendental argument here—that our being effective agents, no less than our being reliable perceivers, is *a necessary presupposition of* our being in the framework of epistemic evaluation, and so of our being thinking and perceiving beings at all—this expectation is belied by the fact that our being effective agents is the *premise*, not the conclusion of the argument Sellars offers in MGEC: ¶¶80–83, and our "being in the framework"—and thus our being reliable perceivers—is not presupposed by the argument but is instead the necessary condition of our effective agency.

[18] My insertion stands in for Sellars's "such a way": in the previous paragraph (¶66), he notes that what he has been searching for is a way in which it could be *independently* reasonable to accept three epistemic principles (notwithstanding that *one* reason for accepting them is that they belong to an empirically well-confirmed "theory of persons as representers of themselves-in-the-world": ¶40). One of the three principles is that my ostensible perceptions (that I lack grounds for doubting) are likely to be true (¶58).

> account [that] might well be called 'Epistemic Evaluation as Vindication'. Its central theme would be that achieving a certain end or goal can be (deductively) shown to require a certain integrated system of means. . . . [T]he end can be characterized as that of being in a *general* position, so far as in us lies, to *act*, i.e., to bring about changes in ourselves and our environment in order to realize *specific* purposes or intentions. (MGEC: ¶¶67–68).[19]

The basis of this argument is not an incontrovertible claim but a practical commitment. We find ourselves with the goal of, as I'll put it, *effective agency*: of reliably achieving specific purposes by changing ourselves and our environments. But since the changes we make to ourselves and our environments won't reliably achieve our purposes if we cannot reliably perceive how things stand with them, we have strong practical reasons to accept as a basis for inference and action that our perceptual judgments are likely to be true, since this is a necessary condition of realizing our end.[20]

---

[19] Sellars notes that this is an application of a strategy he earlier developed in "two essays on the [reasonableness] of accepting inductive hypotheses" (MGEC: ¶67): IV & NDL. The basic idea for which Sellars argues there is that sound inductive arguments should be interpreted, not as non-deductive arguments that assert their conclusions merely as probable (e.g.: *Black clouds are gathering; so, (probably) it will shortly rain*), but rather as deductively valid *practical* arguments from a legitimate goal of the epistemic community—say, that of "possessing a stock of propositions which are either true or within which a rationally controlled proportion are false" (NDL: ¶43)—to the conclusion that a particular proposition (e.g., that it will shortly rain) is reasonable to accept (in virtue of having a property that shows accepting it to be conducive to the satisfaction of that goal). Indeed, Sellars had already extended this strategy to propose a pragmatic basis for the epistemic authority of first principles in a 1965 manuscript, though the details of this basis were left quite sketchy there. (This is OAFP: see ¶¶13 & 19 for schemata of pragmatic arguments that are structurally parallel to those that Sellars proposed enlisting to explain first principles' authority. The Wilfrid S. Sellars Papers archive at the University of Pittsburgh dates the manuscript to 1965: <https://digital.library.pitt.edu/islandora/object/pitt:31735062220300>. It remained unpublished till 1988.)

[20] The distinction I'm drawing between transcendental and pragmatic arguments, then, is that a transcendental argument begins from an uncontroversial feature of our mental lives and argues that the conclusion is a necessary condition of the possibility of that feature, while a pragmatic argument begins from a reasonable goal we have and shows that accepting the conclusion is a necessary means to realizing that goal. These two types of argument seem to me fundamentally distinct, even if we can sometimes formulate viable arguments of both types concerning a given phenomenon. For instance, discussing "the relation between pragmatism and Kantian-styled transcendental argumentation," Pihlström suggests that Rescher's argument that we must postulate things in themselves in order "for us to be able to make sense of our experience" is "at the same time a pragmatic and a transcendental argument" (2023: 67, 73). But it seems better to me to say that we can form either transcendental or pragmatic arguments in favor of this postulation: if realism is indeed "simultaneously . . . a transcendentally necessary condition for the possibility of certain purposive human activities and . . . a pragmatically useful postulate enabling us to engage in those activities efficaciously" (ibid.: 67), then we can argue *either* that realism is true, given that we unquestionably do engage in the relevant activities, *or* that we pragmatically ought to accept realism, given that it is reasonable for us to aim to engage in these activities

At a first pass, then, Sellars's second anti-skeptical argument is this:

The Pragmatic Argument

1. I have the end of being an effective agent.
2. One can be an effective agent only if one's perceptual judgments are likely to be true.
3. So, it is reasonable for me to accept that my perceptual judgments are likely to be true.

But I'll suggest in §4, in which we assess this argument's prospects, that it requires significant revision if it's to succeed even on its own terms.

Do these two arguments of Sellars's stand or fall together? No. One could accept Sellars's transcendental argument while rejecting the pragmatic argument (whether as unsound or simply as non-responsive to the skeptic): as we'll see in §4, some remarks of Michael Williams's (2009: 173–74) suggest that he'd see matters this way. Conversely, one could accept the pragmatic argument as a sufficient response to the skeptic while thinking that the transcendental argument fails. I take this latter view, and I'll argue for it in the paper's remaining sections, beginning with an assessment of Sellars's transcendental argument—and, indeed, of the general prospects for transcendental arguments against skepticism about empirical justification rooted in semantic externalism.

3. Why Sellars's Transcendental Argument Probably Won't Work

Recall our formulation of Sellars's transcendental argument:

The Transcendental Argument

A. My thoughts and assertions are empirically contentful.
B. My thoughts and assertions can be empirically contentful only if my perceptual judgments are likely to be true.
C. Thus, my perceptual judgments are likely to be true.

---

efficaciously. But the possibility that both argumentative strategies might be viable doesn't establish that they're really of a single type.

This argument is valid, so its evaluation comes down to the question of whether and how we can be justified in accepting its premises. Let's take them in inverse order.

We've already considered Sellars's defense of (B) exegetically in §2. So here, I only want to reiterate that this defense rests on Sellars's acceptance of *semantic externalism*: the thesis that some propositional contents of a person's thoughts and statements depend for their individuation on features of her environment external to her body. On Sellars's view, the meanings of our perceptual terms and concepts are fixed in large part by their role in language-entry transitions—which is to say, by their reliable tendency to be evoked, *ceteris paribus*, (just) by the worldly objects and properties they denote. This is the reason, moreover, that he affirms premise (B) and so holds that we can reject doubts about the likelihood of our perceptual judgments' being true: they wouldn't have the meanings they do, and perhaps wouldn't be meaningful at all, if they didn't generally succeed in correlating with the worldly states of affairs they describe.

What about (A): how can we know that our thoughts are meaningful and what, specifically, they mean? For the modern, Cartesian philosophical tradition, there's no problem about that: we have *privileged access to our mental contents*. And actually, his status as a leading critic of this tradition on this topic notwithstanding, Sellars does agree that we have such privileged access, at least in the modest sense in which I'll construe the thesis: he affirms that a subject can know the contents of her own thoughts non-observationally (i.e. without inferring this knowledge from any perceptual belief).[21] It's admittedly true that EPM's concluding sections are devoted to critiquing the traditional Cartesian account of self-

[21] Note that, though I'm stipulating that non-observational self-knowledge is not arrived at through explicit inference from perceptual beliefs, I'm leaving it open whether the warrant it involves might depend in some more indirect way on some perceptual warrant(s). I'll suggest in a moment that, on Sellars's account of privileged access, this does seem to be the case.

knowledge. Sellars denies that a human being is infallible regarding her own mental states. He even denies that she's authoritative by nature concerning their contents—that privileged access is a primitive phenomenon. Rather, he maintains, our capacity to report our mental states non-observationally is a contingent matter, the result of the institution of a social practice. On his view, then, "the fact that each of us has a privileged access to his thoughts . . . is *built on* and *presupposes* [the] intersubjective status" of concepts pertaining to mental states (EPM: §XV/¶59). Still, this quotation demonstrates that Sellars does regard privileged access (of the modest sort formulated above) as genuine—as a *fact*. He accepts that each of us can have (fallible) non-observational knowledge of her own mental states and contents.

Sellars offers plausible justifications of (A) and (B), then. But are these justifications useful in the context of a response to Humean skepticism about empirical justification? This is less clear.[22] Remember that the chief argumentative burden of offering such a response is avoiding circularity: Sellars's task is to identify a warrant to believe that our perceptual beliefs are likely to be true that doesn't depend on any other perceptual warrants. Call such a

[22] Besides those I go on to note in the text, one reason someone might doubt this is that semantic responses to skepticism have typically been deployed in response to Cartesian skeptical worries like the brain-in-a-vat scenario (Putnam 1981: ch. 1), and one might doubt whether they can extend to the Humean problem that concerns Sellars. But I don't share this particular doubt about Sellars's transcendental argument. It's true that some responses to Cartesian skepticism don't extend to Humean skepticism. (For example, even if the *relevant alternatives* response [Dretske 1970, Stine 1976] succeeds in showing that we typically needn't be in a position to rule out our being brains-in-vats to know that we have hands, this doesn't yield any obvious lesson about how to avoid the vicious circularity in the chain of epistemic dependence on which the Humean problem centers.) But in principle, semantic externalism could found a viable response to skeptical problems of either sort. In particular, if I could know independently of any perceptual warrant both that my thought and talk are empirically meaningful and that (given semantic externalism) this entails that my perceptual beliefs are generally true, this would not merely warrant me in denying that I am in a Cartesian scenario where virtually all my perceptual beliefs are false: it would further represent a non-circular warrant for the likely truth of my perceptual beliefs that would answer Humean skepticism. (Incidentally, Sellars's pragmatic argument, too, is pitched as a response to Humean skepticism but equally addresses Cartesian worries: my *a priori* pragmatic warrant to accept that my perceptual beliefs are likely to be true, and so in turn that skepticism about empirical justification is false, equally extends to an *a priori* pragmatic warrant to accept that I'm not in any scenario in which those beliefs would virtually all be false.)

warrant *a priori*—albeit clearly in a very weak sense.[23] Are Sellars's justifications of (A) and (B) *a priori* in this sense? Not obviously. To take the premises again in inverse order, Sellars's argument for (B) involves not simply an analysis of the idea of empirical meaning or a series of thought experiments, but also, crucially, a description of the process of training in proper use of linguistic expressions. This description is a series of empirical claims, and perception plays a central role in warranting them.

And it's even less plausible that his defense of (A) proceeds *a priori*. After all, Sellars maintains that—no less than in the case of perception—the justification of my particular *introspective* judgments rests on my justification for the general belief that my introspective judgments are likely to be true (MGEC: ¶¶44–46 & 59–66). And in EPM's concluding sections, Sellars strikingly suggests that this general belief is a *contingent* truth: it "turns out—need it have?—that [an individual human being] can be trained to give reasonably reliable self-descriptions, using the language of the theory [of overt behavior as the expression of thoughts], without having to observe his overt behavior" (EPM: §XV/¶59). My justification to believe that my introspective judgments are likely to be true would seem[24] to rest, then, on my justification to believe that I have undergone a social process of training in reliable introspective reporting—and this latter justification is empirical, not independent of perceptual warrants.

---

[23] This weak sense is unusual: it implies that—provided it isn't epistemically dependent on any perceptual warrants—introspective knowledge is *a priori*! I use this weak sense for two reasons: first, because it's the sense relevant to whether Sellars can avoid the Humean charge of rule-circularity; and second, because it's the sense typically used in the McKinsey paradox literature, which will figure heavily in the discussion of the more ambitious construal of Sellars's argument to be discussed later in this section. (On the other hand, note that it is insufficient for a warrant to be *a priori* in even this minimal sense that it doesn't depend on any *particular* perceptual warrant: if it is epistemically dependent even on the subject's having *some perceptual warrant or other*, then it still won't count as *a priori* in the sense employed in this section.)

[24] Why hedge the point like this? Because actually, as noted above, Sellars frames his pragmatic argument as covering introspective (and memory) judgments rather than only perceptual ones, which allows him to maintain that we have an additional, epistemically basic and non-empirical—albeit pragmatic—warrant for our belief that our introspective judgments are likely to be true.

While Sellars argues that we can know (A) and (B), then, he doesn't argue that we can do this *a priori*. But then he hasn't enabled his transcendental argument to perform the task he enlisted it to perform in his epistemological writings: namely, to describe "a way in which it could be *independently* reasonable to accept [that our perceptual judgments are likely to be true,] in spite of the fact that *a* ground for accepting [this] is the fact" that this claim is an element of our general theory of the world, which we regard as empirically well-confirmed (MGEC: ¶66; cf. SK III: ¶44–45). This empirical warrant is our reason for accepting the argument's premises, and so the warrant the argument confers on its conclusion turns out not to be suitably independent. Perhaps that explains why, in other contexts—particularly those bound up with his commentary on Kant—Sellars frames the argument's purport more minimally: "A transcendental argument does not prove that there *is* empirical knowledge—what premises could such an argument have? . . . It simply explicates the concepts of *empirical knowledge* and *object of empirical knowledge*" (TTC: ¶53; cf. KTE: ¶¶11, 45). But a mere explication of the concept of empirical knowledge hardly constitutes a response to Humean skepticism about empirical justification, either.

Still, Sellars's question here is an obviously pressing one: what premises *could* an argument that our empirical judgments constitute knowledge—or, at least, that we're rationally entitled to regard our perceptual judgments as likely to be true—have? In §4, I'll argue that this is a comparative argument of Sellars's pragmatic argument over his transcendental one: it's quite plausible that we can know its premises (or actually, it will turn out, corresponding premises in a suitably modified version of the argument) *a priori*. First, though, in the remainder of this section, I want to forestall an attempt someone might make to salvage the transcendental argument: they might wonder whether it's possible to hit upon *a priori* warrants for (A) and (B) after all. And I want to suggest that the attempt to identify

such warrants and to use them to justify the transcendental argument's conclusion won't ultimately yield a satisfying response to the skeptic—and would require some very implausible bullet-biting along the way.

Suppose we were willing to deny that my warrant for my introspective belief that I'm thinking that (e.g.) water is wet in any way epistemically depends on any perceptual warrant of mine. (Even if I can acquire warrant for such beliefs only once I undergo a process of social training in reliable introspective reporting, we might hazard, once this training succeeds, my warrant for particular introspective beliefs doesn't depend on my warrant to believe that the training took place—or perhaps even on my warrant to believe that my introspective beliefs are reliable in general.) And suppose we were further willing to maintain that that semantic externalism, too, is knowable *a priori*—through analysis of our concept of empirical meaning, thought experiments like Twin Earth, or whatever—and that the particular entailments between our beliefs and worldly states of affairs that follow from it are conceptual in nature, and so are themselves knowable *a priori*. If we were willing to make those claims and could manage to motivate them reasonably well, would we have rendered Sellars's transcendental argument cogent? Would we have a promising response to skepticism about empirical justification on our hands?

To me, anyway, the transcendental argument would still seem unacceptable as a response to skepticism. This is for the obvious reason that, if we construe its premises as *a priori*, the argument would suggest that we can have non-observational knowledge—or at least non-observational justified belief—concerning the external world: we could be justified in believing that particular states of affairs obtain in the natural world simply by introspecting about the contents of our thoughts and then engaging in *a priori* reasoning. The idea that this is possible strikes me, frankly, as absurd. And for what it's worth, my reaction

to this suggestion is pretty widely shared. In fact, it seemed so obvious to many philosophers researching these topics as to yield a paradox—one initially formulated by Michael McKinsey, and thus typically called *the McKinsey paradox*[25]—about whether semantic externalism and privileged access are really compatible. Each seems plausible, and yet, if their conjunction opens the door to *a priori* knowledge of the external world, then they simply can't both be true!

Now, leaving the matter there would mean resting content with a stalemate. After all, as defenders of privileged access to the external world have noted, semantic externalism and (at least a modest form of) privileged access are both extremely plausible. Rather than incurring the intuitive cost of denying them, wouldn't we be better off maintaining them and, by embracing in consequence the possibility of non-observational knowledge of our physical environment, gaining a response to skepticism about empirical justification to boot?[26] My answer is that responses are available to the McKinsey paradox that rule out *a priori* knowledge of the external world without rejecting semantic externalism or privileged access altogether, and so can avoid the latter counterintuitive consequences without incurring the former. Let me briefly sketch three such responses.

First and most straightforwardly, we might follow McKinsey himself in accepting only a qualified version of privileged access that doesn't entail that I can know *a priori* that I'm thinking that water is wet—one on which "one does not have privileged access to one's

---

[25] See McKinsey 1991 and, for a useful summary of the literature on the paradox, Kallestrup 2011.

[26] See Warfield 1998: 138. This isn't the only defense that proponents of *a priori* knowledge of the external world give of their position (though I think it's the strongest). For instance, Sawyer (1998: 532) suggests that the strangeness of such knowledge is dispelled by the fact that—given that the causal interactions and perceptual encounters necessary for the grasp of wide contents will also enable perceptual knowledge of the relevant environmental facts—"introspection will yield knowledge *only* of those empirical facts that the subject could already have come to know via empirical means" (Sawyer 1998: 532). With Brown (2004: 238–39) and Brueckner (2010: 236), I don't agree that the ability to come to know environmental facts *a priori* is unproblematic as long as we first knew them empirically.

possession of logically wide mental properties" (2002: 204), that is, properties that conceptually imply the existence of contingent external objects. Second, we might deny with Anthony Brueckner (2010: chs. 5 & 22) that I can know *a priori* which worldly states of affairs are implied by the contents of my thoughts, even given semantic externalism. Compare: nothing follows about my external environment from my *a priori* knowledge that I'm thinking that phlogiston is plentiful around me, even given semantic externalism, since <phlogiston>, though it purports to denote a natural kind, fails actually to do so. Similarly, I can't infer *a priori* from the fact that I'm thinking that water is wet to a conclusion about my external environment unless I know *a priori* that my concept <water> isn't empty but successfully denotes a natural kind—and so unless I have *independent* warrant *a priori* to accept that the relevant worldly states of affairs obtain. Brueckner assumes that I can't have this, but, third, Crispin Wright (2003: 68–69) argues that we're *a priori entitled* to assume that concepts we must use succeed in referring. Even so, however, this doesn't enable me to acquire *a priori* knowledge of the relevant states of affairs from my knowledge of my <water>-thoughts and of their conceptual implication of those states of affairs. For this argument would presuppose my *a priori* warrant to accept that those states of affairs obtain in purporting to provide me with the same, and so would exhibit *failure of warrant-transmission* (ibid.: 69).[27]

I won't defend any of these solutions as against the others, but in my view, their plausibility strongly suggests that the proponent of *a priori* knowledge of the external world is

[27] Warrant-transmission fails when one acquires warrant to believe the premises of a deductively valid argument without thereby acquiring warrant to believe its conclusion. One key reason this can occur is because one was justified in accepting the premises only given one's prior warrant to accept the conclusion. (Wright gives this example: 'Jones has scored a goal' entails 'Jones is playing football'. But even if one knows this, one can come to know the former claim without acquiring any warrant to accept the latter. For if one had lacked prior warrant to believe the latter, one wouldn't have been warranted in believing the former on the basis of seeing Jones kick the ball into the net.) See Wright 2002: §§II–III.

mistaken when she suggests that only admitting such knowledge will allow us to affirm both modest privileged access and semantic externalism. Adequately motivating a transcendental argument like Sellars's as a response to the skeptic, then, would require not only establishing that self-knowledge concerning mental contents and knowledge of semantic externalism are *a priori*, but further arguing that the three solutions just mentioned all fail—and arguing this with enough force to overcome the strong intuitive implausibility of privileged access to the external world. I doubt very much whether that can be done.

In any event, we should note finally that a transcendental argument of this sort would have only very limited force against skeptical hypotheses even if it were cogent. For it's consistent with semantic externalism that my <water>-thoughts should refer to H_2O even if I was transported while asleep last night to Twin Earth (which contains no H_2O but only the superficially-indistinguishable compound XYZ), or that my <dog>-judgments should refer to dogs even if all the dogs in my local environment mysteriously vanished yesterday and were replaced by perfect simulacra. All that externalism requires for these terms to maintain their references is that the subject has a history of extensive interaction with entities of the right kind in the not-too-distant past. But then it provides no resources for refuting the skeptical worry that, though I had reliably accurate beliefs about my external environment for most of my life, my beliefs about its current features are radically mistaken. This skeptical worry is no less pressing than the worry that I've always been radically mistaken about my environment, so we still need a solution to it. And if one can be found, it may well have the resources to answer the broader worry as well, rendering Sellars's transcendental argument superfluous even in that context. [28]

[28] This argument obviously draws on the switching scenarios discussed in Burge 1988 and Boghossian 1989, as well as the familiar "recent envatment" scenario commonly deployed to mitigate the utility of Putnam's (1981:

Therefore, while the transcendental argument represents Sellars's most direct response to skepticism, I don't think it's a very promising one. Let's turn now to consider whether his pragmatic argument, while a more concessive response, may yet fare better.

4. The Value of Sellars's Pragmatic Argument

Recall our formulation of Sellars's pragmatic argument:

> *The Pragmatic Argument*
>
> 1. I have the end of being an effective agent.
> 2. One can be an effective agent only if one's perceptual judgments are likely to be true.
> 3. So, it is reasonable for me to accept that my perceptual judgments are likely to be true.

This argument has some initial plausibility—or, at least, its premises do. (1) seems hard for any of us to deny. Rational agency is, as Crispin Wright remarks, "nothing we can opt out of" (2004: 198): we cannot but act, and in acting we aim to be effective in realizing our intentions. (2) is similarly plausible. As Sellars notes, generally true perceptual judgments amount to a generally reliable cognitive map of one's environment, and without such a reliable map to go by, we would be too likely to find ourselves bumping into things or getting devoured by unobserved predators to reliably achieve our ends.

The primary problem facing the argument is that its validity is suspect, and that on two grounds. First, (3) doesn't follow from (1) and (2) without an auxiliary premise—and, indeed, a rather dubious one:

> AP. It's reasonable for me to accept that any necessary condition for the realization of my ends will obtain.

---

ch. 1) anti-skeptical argument from externalism: see, e.g., Smith (1984). (Regarding Twin Earth, see Putnam 1975.)

I'll construe *acceptance* in Sellars's pragmatic argument as a disposition to assume a claim as a premise in inference and as a basis for action.[29] Thus interpreted, (AP) seems clearly false: there might be some state of affairs whose realization would be valuable enough to warrant my adopting its realization as an end (i.e. my trying to realize it), and yet whose realization remains unlikely to the point that my acting straightaway as if all the necessary conditions of its realization will obtain (i.e. as if it will indeed be realized) nevertheless remains unreasonable. (Consider: there are possible lotteries where the odds of success are sufficiently high relative to the payout and cost of entry to render trying to win reasonable, but sufficiently low absolutely that acting in advance as though one will win remains patently unreasonable.)

Second, (1) is a psychological, descriptive premise: it describes an end we have—and, seemingly, cannot help having. But (3) is a normative conclusion, and (1) doesn't seem capable of founding an argument for it. After all, the mere fact that I adopt some end doesn't seem to *justify* me in anything much: it doesn't justify me in accepting that it will be realized, or even in taking necessary means to it.[30] If the end is itself an unreasonable one, then all that would be reasonable for me to do, having adopted it, is to reverse course and renounce it. If the argument is to yield (3), then, it needs to begin from a normative premise.

---

[29] Here I follow Elgin's (2010: 64) account of acceptance. (Granted, this isn't how Sellars himself uses the term 'accept.' He "use[s] 'accept', in the first instance, as roughly equivalent to 'come to believe'" [NDL: ¶28n2], and he notes that claims can be reasonable to accept in this sense without being reasonable to use as a basis for action [IV: §XVIII]. But since belief seems non-voluntary while Elginian acceptance seems voluntary, and since pragmatic reasons are arguably simply the "wrong kind of reason" for belief—see Hieronymi 2005—but arguably appropriate reasons for Elginian acceptance, I think the latter, not the former, construal of acceptance yields the most defensible version of Sellars's pragmatic argument.)

[30] Does the principle that 'ought' implies 'can' undermine this objection, jointly with the concession that the end of effective agency is psychologically necessary? No. That I cannot but adopt the end of effective agency might *refute* the claim that I *ought not* take necessary means to realizing it, but it certainly doesn't *entail* the claim that it's *positively reasonable* for me to do so.

These objections require revisions to Sellars's pragmatic argument if it's to succeed. In other work (Klemick 2024b), I motivate the objections and the particular revisions I propose at length. Here, though, I'll just state and briefly explain this revised version:

The Modified Pragmatic Argument

1*. I have the reasonable end of being an effective agent.

2*. One can be an effective agent only if one accepts that one's perceptual judgments are likely to be true.

4. It is reasonable for one to take the necessary means to one's reasonable ends.

3. So, it is reasonable for me to accept that my perceptual judgments are likely to be true.

Swapping (1) for (1*) solves the second problem just considered for the original argument: the argument now begins from a normative premise, not a descriptive one. Moreover, this premise seems no less plausible than (1). After all, it's valuable to have control over our empirical circumstances. If we lacked it, we would be at the mercy of our environments. And not only would that make it likely that great harms would befall us, but those harms would be even worse for us because of our total powerlessness to prevent them. So, control over our empirical circumstances is a great good. But only by acting, and by striving to act effectively, can this great good be attained. And that makes the end of effective agency reasonable.

And swapping (2) for the conjunction of (2*) and (4) solves the first problem considered above.[31] Where (AP) is quite implausible, (4) seems obvious. Further, (2*) is no less defensible than (2), since my reliable cognitive map of my environment won't enable me to be an effective agent simply because I have it—I further must actually be disposed to

[31] Another problem it solves is that, as Sellars himself notes elsewhere, it's necessary for a practical argument to be good—where goodness stands to practical arguments as soundness stands to deductive theoretical ones—that the action its conclusion endorses as reasonable must actually serve to bring about the end its basic premise notes (OAFP: ¶34), and (2*) asserts this of my accepting that my perceptual judgments are likely to be true while (2) did not.

employ this map in my reasoning and action.[32] I thus take the modified pragmatic argument to avoid the two apparent problems for Sellars's original argument. Of course, this hardly establishes conclusively that it succeeds. But I do think it shows that some argument much like Sellars's pragmatic one carries some *prima facie* weight: it's plausible, at least, that it succeeds on its own terms.[33]

One reason I can rest content with this very brief defense of the argument's success on its own terms is that I'm likely preaching to the choir. Sellars commentators—the most natural audience, after all, for a close reading of Sellars's treatment of skepticism!—who have explicitly considered whether to attribute something like the pragmatic argument to him have seemed content to grant that it proves its conclusion, for whatever that's worth. What they've questioned is whether it's worth very much at all. DeVries notes that Sellars offers the pragmatic argument partly out of dissatisfaction with a "this or nothing" response to skepticism about perceptual justification. On this response, we simply have to accept that our perceptual beliefs are likely to be true, since if they aren't, we won't have any justified empirical beliefs (MGEC: ¶¶54–55). But while the pragmatic argument is "stronger" than

---

[32] But must this disposition be premised on a metacognitive acceptance that this map is reliable—that my perceptual beliefs are likely true? Couldn't I simply be disposed to use the map without taking any reflective attitude toward its reliability? Indeed, aren't the other animals like this? I see two possible replies to this thought, and I'm sympathetic to both. First, we might deny that one can be an effective rational agent with genuine control over one's empirical circumstances while lacking any reflective attitude toward the view of the world on which one acts. (Genuine control, we might say, requires a measure of rational *self*-control.) More concessively, second, we might note that it's simply a fact that, unlike the other animals, "the human mind *is* self-conscious in the sense that it is essentially reflective" (Korsgaard 1996: 92), and so that we humans can maintain our disposition to think and act in light of our view of the world only through reflectively committing to doing so.

[33] If so, I think it's further plausible that the argument represents a significant advance on previous pragmatist responses to skepticism (with the exception of some lesser-known elements of Peirce's response to skepticism, which, in my view, anticipate Sellars here; see Klemick 2024a). For those responses have typically hinged, as in James and (better-known strands in) Peirce, on some form of *anti-realism*—relativistic or epistemic accounts of truth, or else idealism or phenomenalism about the material world—or else, as in Hume and (arguably, at least at points) Wittgenstein, on a retreat to *psychologism*: the defense of our basic epistemic commitments, not as rationally warranted, but only as psychologically inevitable for animals like us. Without compromising realism about truth or the material world, the modified Sellarsian argument would secure genuine rational warrants for our trust in perception, and so would constitute a much more attractive pragmatist anti-skeptical stance than its predecessors. (I defend these claims, too, more fully in Klemick 2024b.)

the "this or nothing" response (since it initially seems we can concede that we lack justified empirical beliefs and just go on acting as we do anyway, but we can't simply accept a loss of agency), deVries suggests that it's "not significantly so" (2005: 140). Even more pessimistically, Williams argues (in effect) that we should interpret Sellars as propounding not the pragmatic but only the transcendental argument, on the grounds that the pragmatic argument simply fails to respond to the skeptic about empirical justification. In defense of accepting the likely truth of our perceptual judgments, the pragmatic argument offers only the "*strategic*" reason that "as finite knowers in a world we never made, we have no choice but to sign up to some set of epistemic principles" (2009: 173). It offers no *epistemic* warrant for this acceptance, and so cannot secure our justification to hold—but at most our *non-culpability* in holding—our body of empirical beliefs. In conclusion, then, I'll respond by explaining the problem to which I take the (modified) Sellarsian pragmatic argument to constitute an effective response.

It's true that the argument doesn't meet skepticism about empirical justification head-on: Williams is right that it doesn't offer epistemic reasons for the claim that our perceptual beliefs are likely to be true. Rather, it constitutes a *skeptical solution* to such skepticism in Kripke's (1982: 66) sense,[34] one that concedes that it cannot be refuted but aims to establish that, nevertheless, "our ordinary practice or belief is justified because – contrary appearances notwithstanding – it need not require the justification the skeptic has shown to be untenable." This claim is one that Sellars would naturally have thought it important to establish. For the idea that the rationality of our ordinary practice and belief really *does* hinge on the sort of justification whose tenability the skeptic challenges lay at the

[34] This has been noted by Wright (2004: 206–7), whose response to skepticism is in many respects similar to Sellars's pragmatic argument (Williams 2009: 172n12).

very core of the philosophical system of one of the figures against whom Sellars most frequently defined his emerging views throughout the first decade of his scholarship: his Harvard teacher C. I. Lewis.[35]

The driving force behind Lewis's philosophy was his deeply felt "impatien[ce] of those who seem not to face the sceptical doubt seriously" (1970: 3). The tone of his treatments of the subject is far more personal than most writing in analytic epistemology, as when he concludes a late defense of his views with this final paragraph:

> I consider skepticism something worse than unsatisfactory; I consider it nonsense to hold or to imply that just any empirical judgment is as good as any other—because none is warranted. A theory which implies or allows that consequence is not an explanation of anything but merely an intellectual disaster. (ibid.: 330).

Lewis's phenomenalist foundationalism is now widely rejected.[36] But he held that view because he thought it represented the only possible escape from skepticism, which escape is the philosophical task on which everything hangs—and specifically, on which hangs the rationality of all our practical commitments and undertakings:

> that jejune character of consistent skepticism which Hume himself finally admitted, is indicated by the implication of it for action . . . Without [knowable] 'necessary connections' [between matters of fact] there could be no foreseeable consequences of any active attitude; and without such determinable consequences action could not be genuine—the very idea of it would be empty. The skeptic who does not, like the ancient Cynic, refuse to turn out for a wagon, is only play-acting. Consistently he can take no active attitude; not even the attitude of not taking attitudes. And whoever can thus divest himself of his active nature—and without trying—must arouse our wonder if not excite our admiration. At least he will not take his skepticism seriously, or ask us to, since he takes *nothing* seriously. (Lewis 1946: 228).

If skepticism about empirical justification cannot be answered, Lewis thinks, there's nothing it's reasonable to believe about the world in which we act. But then there's nothing it's

---

[35] On Sellars's intellectual relationship to Lewis, see Sachs 2017: ch. 3, O'Shea 2018 & 2021, and Klemick 2024c.

[36] So much so that, while it was once uncontroversial to attribute this stance to him, recent commentators have mostly denied outright that he held either component view. For recent treatments suggesting that the traditional attribution was on the right track, see O'Shea 2021, Browning 2022, and Klemick 2020.

reasonable to do—and so, in turn, nothing worth caring about, since we can never make our concern concretely matter in any case. Therefore, he concludes, if we're to take *anything* seriously, we must take skepticism seriously: we must try to answer it, to meet the skeptic's challenge head-on.

In his early writings, Sellars persuasively poked many holes in Lewis's own answer to skepticism: his idea that we have semantically and epistemically autonomous observations (so-called "apprehensions of the given") that are capable of satisfactorily founding warrants for our ordinary objective empirical beliefs. But I don't think those early writings offer a positive response to Lewis's challenge: they don't contain either a direct answer to the skeptic or an argument that such an answer isn't required to preserve the rationality of our practical commitments. Now, nobody could plausibly maintain that Sellars was as worried by skepticism as Lewis was—indeed, very few philosophers could match Lewis on that front! Still, this paper has shown that he was significantly more concerned to answer it than many commentators have supposed. But while his transcendental argument represents his most direct attempt to do so, I argued in §3 that it doesn't obviously succeed in that task.

The pragmatic argument is valuable, then, because in effect it constitutes a response to Lewis's challenge—and, in my view, Sellars's most persuasive one. It shows that we can give up on the task of directly answering the skeptic about empirical justification without conceding that we have no more reason to base our inquiry and action on any particular empirical belief rather than any other, and so without conceding that no actions are reasonable and that nothing really merits being taken seriously. For even if we have no independent evidence for our basic empirical epistemic principles—centrally, that our perceptual beliefs are likely to be true—we have decisive practical reasons for accepting this claim as a basis for our thought and action, and so for believing empirical claims that cohere

with strong perceptual evidence, rejecting ones that conflict with such evidence, and acting accordingly.[37] Sure, it would be great if, beyond that, we *did* have a non-circular justification to offer in defense of such principles. But as Wright quips: "Good luck to all philosophers who quest for such a demonstration" (2004: 207). Indeed, we've seen in §3 that, at least with respect to transcendental arguments from semantic externalism, the bar to clear is higher than Sellars and some Sellarsians have assumed. For those, then, who become convinced that the quest for some such direct response to skepticism is quixotic, a version of Sellars's pragmatic skeptical solution may prove invaluable. At least, they may plausibly regard it as the most promising contribution he made in response to the skeptical problem.[38]

---

[37] This strategy has clear Kantian roots, since Kant, too, offers practical defenses of our warrant to assume certain fundamental principles (that cannot be theoretically justified) as regulative for inquiry and/or as necessary postulates of practical reason. (Regarding inquiry, see Appendix to the Transcendental Dialectic in the First Critique; regarding practice, see the Dialectic of Pure Practical Reason in the Second Critique.) In this respect (among others), Sellars follows in the lineage of Peirce's deeply Kantian pragmatism (see Misak 2013: §3.9 & pp. 218–19; and again, compare Sellars's pragmatic argument to the interpretation of Peirce's epistemology in Klemick 2024a). Lewis, too, is an important figure in this lineage, but he partly betrays it through his succumbing to the "phenomenalist temptation" (O'Shea 2016: §II), which reflects his perhaps rather blinkered understanding of Kant's importance primarily as the philosopher who "followed scepticism to its inevitable last stage, and laid his foundations where they could not be disturbed" (Lewis 1970: 3). (O'Shea has previously noted the Kantian character of Sellars's justification of our fundamental epistemic principles in MGEC—see his 2011: §§IV–V—but, in highlighting only the transcendental argument without the pragmatic one, doesn't highlight the pragmatist dimension of its Kantian heritage. He does identify other Kantian pragmatist strains in Sellars, however, in O'Shea 2016: §III.)

[38] Thanks to David James Barnett, Cheryl Misak, Steven Levine, Gurpreet Rattan, Andrew Sepielli, and several anonymous referees for very helpful discussion of versions of this paper. I also gratefully acknowledge support from the National Endowment for the Humanities in the form of a stipend to attend its 2019 summer seminar "Philosophical Responses to Empiricism in Kant, Hegel, and Sellars," which stimulated my development of this paper. (And thanks to all the summer school participants for wonderful conversations!)

Bibliography

§1. Sellars's Works (Cited by Abbreviation)

EPH. 1974. *Essays in Philosophy and Its History*. Dordrecht: Reidel.

EPM. 1956/1997. "Empiricism and the Philosophy of Mind." Monograph reprint. Cambridge, MA: Harvard University Press, 1997.

IV. 1964/1974. "Induction as Vindication." Reprinted in *EPH*, 367–416.

KTE. 1967/2002. "Some Remarks on Kant's Theory of Experience." Reprinted in *KTM*, 269–82.

KTM. 2002. *Kant's Transcendental Metaphysics*. Edited by J. Sicha. Atascadero: Ridgeview.

MGEC. 1979. "More on Givenness and Explanatory Coherence." In G. Pappas (ed.), *Justification and Knowledge*, 169–82. Dordrecht: Reidel.

NDL. 1970/1974. "Are There Non-Deductive Logics?" Reprinted in *EPH*, 417–38.

OAFP. 1988/2002. "On Accepting First Principles." Reprinted in *KTM*, 443–53.

SK. 1975. "The Structure of Knowledge." In H.-N. Castañeda (ed.), *Action, Knowledge and Reality*, 295–347. New York: Bobbs-Merrill.

SM. 1967/1992. *Science and Metaphysics*. Reprint. Atascadero: Ridgeview, 1992.

SRI. 1965/2011. "Scientific Realism or Irenic Instrumentalism." Reprinted in *Philosophical Perspectives: Metaphysics and Epistemology*, 146–72. Atascadero: Ridgeview, 2011.

SRLG. 1954/2007. "Some Reflections on Language Games." Reprinted in *In the Space of Reasons*, ed. K. Scharp & R. B. Brandom, 28–56. Cambridge, MA: Harvard University Press, 2007.

TC. 1962/1991. "Truth and 'Correspondence.'" Reprinted in *Science, Perception, and Reality* (reprint), 197–224. Atascadero: Ridgeview, 1991.

TTC. 1970/2002. "Toward a Theory of the Categories." Reprinted in *KTM*, 321–39.

WSNDL. 2018. *Wilfrid Sellars Notre Dame Lectures 1969–1986*. Revised edition. Ed. D. Landy & P. V. Amaral. Atascadero, CA: Ridgeview.

§2. Other Works

Boghossian, Paul. 1989. "Content and Self-Knowledge." *Philosophical Topics*, 17: 5–26.

Brandom, Robert. 2013. "Global Anti-Representationalism?" In H. Price, *Expressivism, Pragmatism, and Representationalism*, 85–111. Cambridge: Cambridge University Press.

———. 1997. "Study Guide." In EPM, 119–81.

———. 1994. *Making It Explicit*. Cambridge, MA: Harvard University Press.

Brown, Jessica. 2004. *Anti-Individualism and Knowledge*. Cambridge, MA: The MIT Press.

Browning, Jacob. 2022. "C.I. Lewis: Pragmatist or Reductionist?" *International Journal of Philosophical Studies*, 30/2: 109–26.

Brueckner, Anthony. 2010. *Essays on Skepticism*. Oxford: Oxford University Press.

Burge, Tyler. 1988. "Individualism and Self-Knowledge." *The Journal of Philosophy*, 85/11: 649–63.

Davidson, Donald. 2001. *Subjective, Intersubjective, Objective*. Oxford: Oxford University Press.

DeVries, Willem. 2019. "Picturing: Naturalism and the Design of a More Ideal Truth." In P. Giladi (ed.), *Responses to Naturalism*, 230–49. New York: Routledge.

———. 2005. *Wilfrid Sellars*. Montreal/Kingston: McGill-Queens University Press.

DeVries, Willem and Timm Triplett. 2000. *Knowledge, Mind, and the Given*. Indianapolis: Hackett.

Dretske, Fred. 1970. "Epistemic Operators." *The Journal of Philosophy*, 67/24: 1007–23.

Elgin, Catherine. 2010. "Persistent Disagreement." In R. Feldman & T. Warfield (eds.), *Disagreement*, 53–68. Oxford: Oxford University Press.

Hieronymi, Pamela. 2005. "The Wrong Kind of Reason." *The Journal of Philosophy*, 102/9: 437–57.

Kallestrup, Jesper. 2011. "Recent Work on McKinsey's Paradox." *Analysis*, 71/1: 157–71.

Klemick, Griffin. 2024a. "Of Hopes and Hinges: Peirce, Epistemic Constraints on Truth, and the Normative Foundations of Inquiry." *Erkenntnis*. <https://doi.org/10.1007/s10670-024-00854-z>.

———. 2024b. "Perceptual Justification and the Demands of Effective Agency." *Synthese*, 203/34: 1–20.

———. 2024c. "Sellars's Core Critique of C. I. Lewis: Against the Equation of Aboutness with Givenness." *Archiv für Geschichte der Philosophie*, 106/1: 106–36.

———. 2020. "C. I. Lewis was a Foundationalist After All." *History of Philosophy Quarterly*, 37/1: 77–99.

Korsgaard, Christine. 1996. *The Sources of Normativity*. Cambridge, MA: Harvard University Press.

Kripke, Saul. 1982. *Wittgenstein on Rules and Private Language*. Cambridge, MA: Harvard University Press.

Levine, Steven. 2019. *Pragmatism, Objectivity, and Experience*. Cambridge: Cambridge University Press.

Lewis, C. I. 1970. *Collected Papers of Clarence Irving Lewis*. Edited by J. D. Goheen and J. L. Mothershead. Stanford, CA: Stanford University Press.

———. 1946. *An Analysis of Knowledge and Valuation*. La Salle: Open Court.

Matsui, Takaaki. 2021. "Inferentialism and Semantic Externalism: A Neglected Debate between Sellars and Putnam." *British Journal for the History of Philosophy*, 29/1: 126–45.

McDowell, John. 2009. *Having the World in View*. Cambridge, MA: Harvard University Press.

———. 1994/1996. *Mind and World*. Paperback reprint. Cambridge, MA: Harvard University Press, 1996.

McKinsey, Michael. 2002. "Forms of Externalism and Privileged Access." *Philosophical Perspectives*, 16: 199–224.

———. 1991. "Anti-Individualism and Privileged Access." *Analysis*, 51/1: 9–16.

Misak, Cheryl. 2013. *The American Pragmatists*. Oxford: Oxford University Press.

O'Shea, James. 2021. "Relocating the Myth of the Given in Lewis and Sellars." In Q. Kammer, J.-P. Narboux, & H. Wagner (eds.), *C.I. Lewis: The* A Priori *and the Given*, 195–215. New York: Routledge.

———. 2018. "The Analytic Pragmatist Conception of the A Priori: C. I. Lewis and Wilfrid Sellars." In M. Baghramian & S. Marchetti (eds.), *Pragmatism and the European Traditions*, 203–27. New York: Routledge.

———. 2016. "Concepts of Objects as Prescribing Laws: A Kantian and Pragmatist Line of Thought." In G. Gava & R. Stern (eds.), *Pragmatism, Kant, and Transcendental Philosophy*, 196–216. New York: Routledge.

———. 2011. "How to be a Kantian *and* a Naturalist about Human Knowledge." *Journal of Philosophical Research*, 36: 327–59.

———. 2007. *Wilfrid Sellars*. Cambridge: Polity.

Pihlström, Sami. 2023. *Realism, Value, and Transcendental Arguments between Neopragmatism and Analytic Philosophy*. Cham, Switzerland: Springer.

Putnam, Hilary. 1981. *Reason, Truth and History*. Cambridge: Cambridge University Press.

———. 1975. "The Meaning of 'Meaning.'" In K. Gunderson (ed.), *Language, Mind, and Knowledge*, 131–93. Minneapolis: University of Minnesota Press.

Rorty, Richard. 1991. *Objectivity, Relativism, and Truth*. Cambridge: Cambridge University Press.

Rowlands, Mark, Joe Lau, and Max Deutsch. 2020. "Externalism about the Mind." *Stanford Encyclopedia of Philosophy* (Winter 2020 edition), edited by E. N. Zalta. <https://plato.stanford.edu/archives/win2020/entries/content-externalism/>.

Sachs, Carl. 2017. *Intentionality and the Myths of the Given*. Reprint. London: Routledge.

Sawyer, Sarah. 1998. "Privileged Access to the World." *Australasian Journal of Philosophy*, 76/4: 523–33.

Smith, Peter. 1984. "Could We Be Brains in a Vat?" *Canadian Journal of Philosophy*, 14/1: 115–23.

Stine, G. C. 1976. "Skepticism, Relevant Alternatives, and Deductive Closure." *Philosophical Studies*, 29/4: 249–61.

Warfield, Ted. 1998. "A Priori Knowledge of the World: Knowing the World by Knowing Our Minds." *Philosophical Studies*, 92: 127–47.

Williams, Michael. 2009. "The Tortoise and the Serpent." In W. deVries (ed.), *Empiricism, Perceptual Knowledge, Normativity, and Realism*, 147–85. Oxford: Oxford University Press.

Wright, Crispin. 2004. "Warrant for Nothing (and Foundations for Free)?" *Aristotelian Society*, Supplementary Volume 78/1: 167–212.

———. 2003. "Some Reflections on the Acquisition of Warrant by Inference." In S. Nuccetelli (ed.), *New Essays on Semantic Externalism and Self-Knowledge*, 57–77. Cambridge: MIT Press.

———. 2002. "(Anti-)Sceptics Simple and Subtle." *Philosophy and Phenomenological Research*, 65/2: 330–48.