Book Review *Mind* Forthcoming

Critical Notice of Gila Sher's *Epistemic Friction*

Epistemic Friction: An Essay on Knowledge, Truth, and Logic, by Gila Sher. Oxford: Oxford University Press, 2016. Pp. 370.

Terry Horgan, University of Arizona

Gila Sher's *Epistemic Friction* is a bold and ambitious book, with many interesting things to say not only about knowledge, truth, and logic (the three topics mentioned in her subtitle) but also about matters ontological. It often requires the reader (this reader, anyway) to construe it hermeneutically, but repays the effort of doing so.

She coins the expression 'epistemic friction' to refer to *constraints* on a system of knowledge, coming from both the world and the mind. She says, "The world as the object or target of our theories restricts what we can truly say about it, and the mind restricts our theories both voluntarily and involuntarily" (p. 3). Borrowing terminology from Shapiro (1991), she describes her project as *foundation without foundationalism*. "The key idea," she says, "…is that there is no inherent connection between grounding our system of knowledge in reality and doing so in a strictly ordered manner…. Foundation-without-foundationalism is committed to avoiding *vicious* circularity/regress, but since not all non-strictly-ordered grounding is viciously circular/regressive, the road is open to a non-foundationalist yet foundational methodology" (p. 23). Here I will focus on just a few central elements in her rich and broad-ranging articulation and defense of this epistemological approach.

1. Sher is very critical of epistemological empiricism, as she construes this position. She interprets Quine as a (radical) empiricist, and opposes him in this regard. But although she regards Quine's center/periphery model of knowledge as too empiricist and too static, she does embrace a variant of it—a dynamic variant. So, how is she construing empiricism, and how does this connect to her dynamic center/periphery model of knowledge? In the preface she says the following, about Quine's model of knowledge.

> Quine's model has a significant interface with both reality and mind (periphery and center); it is a rich holistic model, with an elaborate network of connections between diverse units of

knowledge; and it rejects the traditional divisions of units of knowledge into those grounded in reality and those grounded *solely* in the mind. But while I adopt these elements of Quine's model, I renounce others…. [Its] periphery is limited to observational sentences and its center to pragmatic considerations and conventional postulations…. There is no room in Quine's model for a veridical challenge to logic and mathematics as branches of knowledge in their own right.

This limited view of abstract knowledge is closely connected to Quine's deep, and quite radical, empiricism—another target of my criticism. Not only is Quine's conception of center and periphery relatively narrow, but so is his conception of mind and world. Quine never asks whether objects in the world have abstract features, nor does he ask whether humans have resources for cognizing such features. To me, it is inconceivable that humans would reach the level of knowledge they have without a significant contribution of intellect (as something distinct both from sensory perception and from pragmatic conventions), but Quine completely neglects the role of intellect in knowledge. He considers, and rightly rejects, supernatural means of discovery and justification, but human intellect is not supernatural. Human intellect cannot be identified with either telepathy or clairvoyance, nor is it related to Greek deities, or the like. In the entire Quinean corpus there is no consideration of intellect as a crucially, or even potentially, significant cognitive corpus…. (pp. ix-x)

I myself sympathize with the spirit of these remarks. Creative intellect contributes very significantly to the acquisition of knowledge about the world, and the "friction" from the mind-independent world that we humans encounter when seeking to know about the world is hardly limited to the immediate deliverances of sensory-perceptual experience. For one thing, knowledge-generation is a richly *abductive* matter, and the concepts that get invented and deployed in the course of abductive theorizing in science draw heavily on the creative work of the human intellect. Explanatory understanding of the world is a much richer phenomenon than simply being able to predict upcoming sensory stimulations; and explanatory understanding requires very rich input from intellect and its capacities.

But in what sense, if at all, does Quine completely neglect the role of intellect, as something distinct from pragmatic conventions? I take her to mean that for Quine, the reifications of intellect do not really contribute anything to human knowledge beyond successful predictions of potential upcoming observations. In a text that she herself does not cite or quote, Quine wrote:

> [I]f we transform the range of objects of our science in any one-to-one fashion, by reinterpreting our terms and predicates as applying to new objects instead of the old ones, the entire evidential support of our science would remain undisturbed…. The conclusion is that there can be no evidence for one ontology as over against another, so long anyway as we can express a one-to-one correlation between them…. Certainly we are dependent on a familiar ontology of middle-sized bodies for the *inception* of reification, on the part both of the individual and of the race; but once we have an ontology, we can change it with impunity…. The very notion of object, or of one and many, is indeed as parochially human as the parts of speech; to ask what reality is *really* like, however, apart from human categories, is self-stultifying. It is like asking how long the Nile really is, apart from parochial matters of miles of meters. Positivists were right in branding such metaphysics as meaningless. (Quine 1992, pp. 8-9)

As against this radically empiricist conception of both the meaning of the sentences that collectively constitute a body of knowledge and of the nature of that knowledge itself, I take Sher to be claiming that some posits of human intellect figure in sentences that correspond to what reality *is* really like, whereas other posits of human intellect do not thus correspond to reality—even if those other posits can be deployed in some body of sentences that is empirically equivalent to the body of knowledge-constituting sentences. In this I applaud her.

When engaging in epistemic projects like seeking explanatory understanding, quite a lot other than Quinean observation sentences and their contents can move to the periphery of a center/periphery structure of knowledge—something that Sher emphasizes by offering us a dynamic variant of Quine's own center/periphery model. As she puts it,

> My own model…is a dynamic model: center and periphery are job descriptions rather than fixed locations, and each discipline moves from periphery to center and vice versa according to the task at hand. (p. x)

This too seems very plausible. Consider, for instance, theorizing in cosmology about the large-scale spatio-temporal structure of the universe. One's abductively best theoretical account here might well turn out to be one according to which, say, the cosmos is finite in spatial extent and yet unbounded—with the large-scale geometry, say, of the three-dimensional surface of a temporally expanding four-dimensional hyper-sphere. One way to dynamically bring such a theoretical hypothesis to the periphery of

our knowledge-system is to point out, as a putative mark against it, that we humans are not capable of literally imagining—literally forming a mental image of—a four-dimensional hypersphere. But then other considerations can be brought to bear at this current periphery, pulling in the other direction and arguably outweighing the evidential significance of this unimaginability phenomenon—for instance, the fact that higher-dimensional generalizations of non-Euclidean geometry are mathematically natural and straightforward, and the fact that if we try to understand the large-scale geometrical structure of the cosmos in a purely Euclidean way we run smack into Kantian antinomy.

Or consider, for another example, debates in philosophy about what kind of logic and semantics is best for accommodating vagueness—which, among other things, requires finding a way to block the infamous sorites paradox. One way to dynamically bring competing proposed systems of logic for vagueness to the periphery of our knowledge system is by noting certain commitments of the different systems and then asking how well these commitments conform with ordinary common-sense thought and talk. Take supervaluationism, for example. Concerning a statement like "For every n, if an n-grained pile of sand is a heap then an n-1 grained pile of sand is a heap," a supervalutionist will affirm this statement's classical negation, and will also affirm the following statement, which is logically equivalent (both under classical two-valued semantics and under supervaluationist semantics) to that classical negation: "There exists an n such that an n-grained pile of sand is a heap and an n-1 grained pile of sand is not a heap." The supervaluationist will also say that although this existential claim is true, it has no true instances. Well, one important count against supervaluationism, while we are holding it at the periphery of the knowledge system, is this: claiming that the existential statement is true despite the fact that none of its instances are true goes deeply contrary to what we ordinarily would tale ourselves to *mean* in affirming that existential statement. Other proposed approaches to the logic and semantics of vagueness can perhaps do better in this respect—perhaps without encountering comparably powerful objections themselves.

At the beginning of Chapter 5 Sher further elaborates her dynamic center/periphery model as follows.

Our model is characterized by:

1. *A broad and open-ended conception of reality, neither Platonist nor nominalist, one that affirms both experiential and abstract features of reality—one reality—and regards them as interconnected.*

2. *A view of human intellect as playing a central role in knowledge, both abstract and empirical. Intellect's role is central not just to the conceptual or pragmatic aspects of knowledge, but also, and significantly, to its veridicality. By affirming the central role of intellect in knowledge, however, the model does not affirm apriorism, nor does it identify non-apriorism with empiricism. The model is neither apriorist nor empiricist, regarding intellect and sensory perception as two essential, yet interconnected, elements of knowledge.*
3. *"Basic" realism: A realism which is more robust than most other forms of realism in regarding reality as both the target and ground of all human knowledge (including abstract, e.g., logical and mathematical, knowledge), yet is more flexible than most other forms of realism with respect to the ways in which a theory can be (substantially) connected to reality.* (p. 73)

Here she appears to be using the term 'empirical' and the label 'empiricism', in such a way that the category 'empirical' largely or completely excludes the operation of what she calls 'intellect' in human knowledge. Thus, she also appears to be construing *empiricism* as an epistemological orientation that also largely or completely excludes the operation of intellect (which fits with her remarks in an above-quoted about Quine's radical empiricism). This way of understanding her is further underwritten by the following remarks:

> Knowledge of the world, or of various aspects of the world, is attained, in our model, by a combination of activities, ranging from sensory perception to conceptualization, abstraction, generalization, reflection, combinatorics, analysis, figuring out, model building (scientific and logical), experiment design, mathematical intuition, and so on. In this essay I place these activities **with the exception of sensory perception**, under "intellect." Intellectual activity, in the sense intended here, is thus any non-sensory cognitive activity, i.e., any cognitive activity that does not principally rely on any of our five senses. I also use "reason" as a synonym for "intellect" in this sense. (p. 84, my emphasis)

Her "basic realism," I take it, affirms (as against the above-quoted passage from Quine), both that it does make sense to ask what reality is *really* like, and also that humans can indeed obtain knowledge about what reality is really like. Such knowledge, of course, would not be something "apart from human

categories" (as Quine put it), but instead would deploy certain human categories to frame claims that *correspond* to what reality is really like. (Cf. Section 2 below.)

Her rejection of what she calls "empiricism," I take it, really amounts to a rejection of the kind of verificationism about meaning that leads to the contention that metaphysical questions about what reality is really like are meaningless. Now admittedly, words like 'empirical' and 'empiricist' are terms of art in philosophy (and in science). But although Quine and the positivists embraced a radical empiricism that asserts in effect that *meaning* is exhausted by matters observational (with Quine treating the whole of science as the unit of meaning, rather than the individual sentence), nonetheless the words 'empirical' and 'empiricism' often are used in philosophy much more broadly—'empirical' for something like "not apriori," and 'empiricism' for the view that there is no a priori synthetic knowledge (or, if one rejects the analytic/synthetic distinction, for the view that there is no a priori knowledge at all). Moreover, Sher herself rejects the analytic/synthetic distinction, and is very dubious about the category of the pure a priori. (She holds that neither logic nor mathematics is purely a priori.) In addition, her dynamic version of the center/periphery model allows that components of one's system of knowledge which Quine treated as always at or near the center actually can, and do, sometimes get dynamically brought to the periphery in the course of inquiry—thereby becoming more thoroughly subject to "friction" from reality than in Quine's own center/periphery model of knowledge. In these respects, her position is arguably even more thoroughly "empiricist" than is Quine's—even though hers is neither a radical, verificationist, empiricism about meaning itself, nor a foundationalist empiricism that treats all knowledge as an inferentially strictly-ordered superstructure built upon the foundation of sensory experience.

2. Sher argues that truth is correspondence to reality, and also that this correspondence takes many forms. She says:

> [A] core, universal principle of truth concerns the connection between mind and reality inherent in it. Traditionally, this principle is called "correspondence", but the form it takes in the present theory differs in several ways from its traditional form(s). In particular, our version of correspondence allows it to take multiple forms, including highly intricate forms, in contrast to the single and simple form—copy (picture, mirror) or isomorphism—that most traditional conceptions of this relation demand…. To signal both the similarities and the differences between

traditional correspondence and the one proposed here, I will call the latter "manifold correspondence". As a starting point, we can briefly formulate the principle of manifold correspondence as follows:

(M-COR) *Truth is a matter of a substantial and systematic connection between thought and reality, a connection that has to do both with the way the world is and the way our mind operates. This connection might be quite intricate and take different forms in different fields. The forms it takes depend both on what aspect of reality a given thought targets and on the cognitive resources available to us for reaching it. Abstracting from differences, this connection holds between a given thought and reality when the aspect of reality it targets is, directly or indirectly, yet systematically, as it says it is.* (p. 186)

While I sympathize strongly with these remarks, again I find myself reading them hermeneutically—and more particularly, as conforming with my own approach to truth. (In Chapter 8 she explicitly compares her views about truth to mine (Horgan 2001), and also to those of Aristotle, Frege, Quine, Field, and Yablo.)

First, I construe M-COR in a way that includes an explicitly *ontological* understanding of the distinction between direct correspondence on one hand, and the various kinds of indirect correspondence on the other hand. Only when the pertinent kind of correspondence is direct do one's claims carry genuine "ontological commitment" regarding what reality is really like. I.e., only then do one's claims embody commitments regarding the items (objects, events, properties, tropes—whatever) that belong to the correct ontology. (This of course is a realist use of Quine's expression 'ontological commitment'—a use which he himself presumably would have repudiated as belonging to meaningless metaphysics.) As one might put it, when the contextually operative kind of correspondence is indirect rather than direct, a true statement's *grammatical* "ontological commitments" (by Quineian criteria) need not be *robust* ontological commitments.

Second, I construe M-COR as identifying truth with correct affirmability under contextually operative, and contextually *variable*, semantic-affirmability standards. Generically, truth normally is a joint product of (a) the contextually operative semantic-affirmability standards, and (b) how things are with reality. But direct correspondence and indirect correspondence are distinct species of this genus.

Sometimes, the contextually operative standards work in such a way that that the statements they govern are ontologically committal in the metaphysical-realist sense; these are direct-correspondence semantic standards. But in many contexts the operative semantic standards instead operate differently: although they do impose certain requirements upon reality in order for a given statement to count as semantically correctly affirmable (i.e., as true), they do not require the correct ontology to include items answering to the statement's posits. These are indirect-correspondence semantic standards.

As a plausible example of a statement that I claim would normally be governed by indirect-correspondence standards of semantically correct affirmability (i.e., of truth), consider this: "Arizona has exactly three public universities." We all know well enough what kinds of evidence would be required to justify this claim, and such evidence would not suffice to become justified in the contention that the right ontology includes putative entities like states, nations, or public universities. And although the positivists were mistaken in regarding truth as a matter of *epistemic* correctness, they were on to an important fact even so: viz., that often there is a comparatively small gap between epistemic correctness and truth (i.e., between epistemic correctness and contextually operative *semantic* correctness). Such small gap-size, in the context of a given discourse, typically is an indicator that in that context, truth is a matter of indirect correspondence.

On Sher's "minimal realism," truth is indeed correspondence to reality; thus—contra contra Quine—true statements really do say "how reality really is." As I understand M-COR, however, when the contextually operative kind of correspondence is indirect rather than direct, true statements say how reality is in such a way that the posits they employ need not necessarily pick out denizens of the correct ontology; i.e., their grammatical ontological commitments need not be robust ontological commitments. Of the innumerable ways reality might be, the actual way-reality-is is a *There-are-exactly-three-public-universities-in-Arizona* way—even though the right ontology need not contain such putative items as public universities, nations, or the State of Arizona. (The grammatical ontological commitments of the statement that Arizona has exactly three public universities are not robust ontological commitments.) And, again contra Quine, once our thought and discourse employ posits like public universities and the State of Arizona, it is not the case that we can change them "with impunity" so long as the class of entailed observation sentences remains unaltered; for, such impunity would go contrary to the indirect-

correspondence affirmability standards that are *actually operative* in prosaic thought and prosaic discourse about universities and States of the Union.

Sher advocates substantivism about truth, as against deflationism—an attitude with which I myself wholeheartedly concur. She says:

> Substantivists (advocates of a substantive theory of truth) differ from deflationists on multiple points: Where deflationists say that "truth is entirely captured by the…triviality…that each proposition specifies its own condition for being true", substantivists say that it is far from being fully captured by this triviality; where deflationists say that "the truth predicate exists solely for the sake of a certain logical need" (Horwich [1990/8]: 2), substantivists say that it exists for other needs as well; where deflationists say that truth is not a deep notion, substantivists say it is; and where deflationists say that a theory of truth cannot be, or need not be, genuinely explanatory, substantivists say it can and should be. (p. 132)

But just what are the other needs served by the notion of truth, and just how can claims about the truth or falsity of first-order claims be explanatory in ways that differ from the explanatory potential of the first-order claims themselves?

Once again I find myself reading her hermeneutically. As I understand M-COR, it is a construal of truth that allows us to fend off quietism about substantive metaphysical debates in philosophy. Deflationists tend to go minimalist across the board about a host of conceptually interconnected notions: truth, property/relation, fact, correspondence. They also tend to think that there is no real subject-matter-neutral notion of existence. So they are apt to say that there is no substantive issue whether numbers exist, since (for example) it follows from the Peano axioms that there exists a natural number greater than 2. And they are apt to say that there's no substantive question whether moral properties or moral facts exist, since (for example) murder is wrong—from which it follows that murder instantiates the property wrongness, that murder's being wrong is a fact, and that the statement 'Murder is wrong' corresponds to that fact.

Crucial to M-COR, as I interpret it, is the contextual variability of semantic correct-affirmability standards. On one hand, this feature allows the approach to accommodate what is right in deflationist or minimalist approaches: the truth-predicate *often* is governed by contextually operative semantic standards under which truth-talk runs smoothly, disquotationally, in tandem with first-order talk; likewise, the

semantically operative semantic standards governing property-talk, fact-talk, correspondence-talk, and existence-talk often work in this same minimalist way vis-à-vis first-order talk. But on the other hand, (a) often those contextually operative semantic correct-affirmability standards will be *indirect correspondence* standards of one kind or another, and (b) the possibility is always there of shifting the contextually operative "score in the language game" (Lewis 1979) into *direct correspondence* standards, thereby posing and addressing substantive metaphysical issues. As the currently fashionable saying goes, the possibility is always there of speaking and thinking in the manner appropriate for being in the "ontology room"—i.e., the manner governed by *direct correspondence* semantic standards, in which the grammatical ontological commitments of one's thought and discourse constitute genuine, robust, ontological commitments. (Indeed, the possibility is always there of shifting the language-game score in such a way that although the newly operative semantic-affirmability standards are still indirect-correspondence standards, nonetheless in pertinent respects they are closer to direct-correspondence standards than before. For instance, we can debate moral realism vs. moral irrealism—in contexts of detached metaethical dispute, as distinct from contexts of engaged first-order moral dispute—while still helping ourselves to posits like universities and the State of Arizona.) So much the worse for metaphysical quietism, and so much the better for metaphysics.

3. My own construal of M-COR leaves room for the possibility that some statements are analytic, at least sometimes and under some contextually operative language-game scores. Indeed, it also leaves room for the possibility that some *existence claims* are analytic—e.g., the statement that there exists a natural number greater than 2. I myself find it very plausible that the truths of pure mathematics are analytic—existence claims included. That is, these mathematical claims are semantically correctly affirmable under the contextual semantic-affirmability standards that *normally* govern thought and talk deploying mathematical posits—and are correctly affirmable solely by virtue of those standards themselves, apart from any "friction" from reality. This view is attractive epistemologically, with respect to understanding the nature of mathematical knowledge. And it is attractive metaphysically, since it is entirely compatible with saying the following after shifting the language-game score into direct-correspondence semantic standards: Numbers and other putative mathematical entities *do not exist*.

Sher, as I noted above, sides with Quine in repudiating analyticity altogether—including in pure mathematics. Against my own approach to mathematical truth, she says that on this view, "[o]rdinal and cardinal terms, like any other mathematical terms, are not connected to reality at all. Instead, they are connected (exclusively) to pure conventions, i.e., in our terminology to the mind" (p. 215). I lack the space to rehearse here her arguments against analyticity and my reasons for finding them unpersuasive. Instead, let me conclude my discussion of her rich and provocative book by briefly replying to the just-quoted charge she levels against me.

The charge seems to assume, in effect, that if the true claims of pure mathematics are analytic, then the mathematical concepts that figure in these claims cannot figure in any claims that are subject to friction from reality—i.e., claims that are subject to being rendered either true or false by how things really are in our world. But that assumption seems to me eminently deniable. Even if *some* claims deploying ordinal and cardinal terms are analytic (including some existence claims), it hardly follows that *all* claims deploying such terms are analytic. The claim that the earth has exactly one moon is empirical; the claim that Arizona has exactly three public universities is empirical; and so on for numerous other claims employing ordinal and cardinal terms.

Imagine a "deductive savant" (I do not say a *mathematical* savant) who is marvelously good at applying rules of deductive logic to certain axioms—e.g., the Peano axioms, the axioms of Euclid—to derive theorems. Suppose this person also confidently mouths the axioms and theorems and confidently embraces them, but has no capacity to *count* and has no idea what you are talking about if you ask, for example, "How many siblings do you have?" Although the axioms and theorems the person mouths and embraces are, in fact, truths of pure mathematics, this person has no capacity at all to deploy mathematical concepts in framing any claims that are subject to friction from reality. Would this person possess *mathematical* concepts, or possess knowledge of analytically true mathematical claims *qua mathematical*? I would say no; and I think it is entirely open to me to say this, consistently with my contention that the truths of pure mathematics are all analytic. Two morals emerge. First, analytic truths in pure mathematics can constrain the correct use of mathematical concepts in claims that are subject to friction from reality, just as purely analytic truths about bachelorhood can constrain the correct use of the concept BACHELOR in claims that are subject to friction from reality—e.g., the claim that Isaac Newton was a lifelong bachelor. (That empirical claim could not be true unless Newton was never married.)

Second, a person utterly unable to appreciate such constraints would not really possess mathematical concepts—however good the person might be at mouthing sentences that happen to express analytic mathematical truths, or at deriving theorems from axioms.

References

Horgan, Terence 2001, 'Contextual Semantics and Metaphysical Realism: Truth as Indirect Correspondence', in M. Lynch (ed.), *The Nature of Truth*. Cambridge, MA: MIT Press, 67-95.

Horwich, Paul, 1990/8, *Truth*. Oxford: Oxford University Press.

Lewis, David 1979, 'Scorekeeping in a Language Game', *Journal of Philosophical Logic* 8: 339-359.

Quine, Willard Van Orman 1992, 'Structure and Nature', *The Journal of Philosophy* 89: 5-9.

Shapiro, Stewart 1991, *Foundationalism without Foundations*. Oxford: Oxford University Press.