

# Science and the Modest Image of Epistemology

*Owen Flanagan*<sup>\*</sup>  
ojf@duke.edu

*Stephen Martin*<sup>\*</sup>  
sjm30@duke.edu

## ABSTRACT

In *Philosophy and the Scientific Image of Man* (1963) Wilfrid Sellars raises a problem for the very possibility of normative epistemology. How can the “scientific image”, which celebrates the causal relation among often imperceptible physical states, make room for justificatory relations among introspectible propositional attitudes? We sketch a naturalistic model of reason and of epistemic decisions that parallels a compatibilist solution to the problem of freedom of action. Not only doesn’t science lead to rejection of our account of normative reasoning, science depends on, sophisticates, and explains how normative reasoning is possible.

The final problem of philosophy is to connect the fact and content of knowledge with its conditions. How does knowing occur in the kind of world that is actually known? Knowing is a fact and must be connected up with the world which the sciences study. Thus a system of philosophy answering this question is the capstone of science.

(Roy Wood Sellars, *Evolutionary Naturalism*, 1922, p. 6)

<sup>\*</sup> Department of Philosophy, Duke University, Durham NC, USA.

## 1. The Space for Reasons in a World of Causes

Does empiricism, the epistemology of the sciences, remain recognizable as epistemology as the sciences advance? Epistemology provides the standards of *justified true belief*, the norms of *warranted assertability*, the criteria for determining the *quality* of evidence and inference. What is the place of such a discipline in the world described and explained in terms of causal interactions among unfamiliar events and processes, some of them imperceptible? (Sellars, 1963).

The problem can be expressed in a variety of ways, all of which are disquieting from the perspective of one who thinks of philosophy, generally, and epistemology, in particular, as foundational, autonomous, antecedent to, or independent of inquiry; or, to put it another way, who thinks of epistemology in a way that makes its core concepts and categories immune from doubts resulting from the very inquiry it promotes and endorses.<sup>1</sup> How does *truth*, specifically the *truth relation*, where ‘p’ is true iff p, fit with the dominant relation that the “scientific image” celebrates, the *causal relation*? How, if they do, do beliefs, which along with *desires* are a fundamental element of folk psychology, survive into the scientific image? How can the sciences preserve the *knowledge relation*, as the perennial philosophy conceives it, if the processes of acquiring beliefs, evaluating beliefs, and consulting epistemic standards are analyzed reductively in terms of the flow of information, where the flow of information is understood in causal-neural-computational terms?

In *Empiricism and the Philosophy of Mind* (1997) Sellars marks critical thinking, the operation of questioning, evaluating, and deciding to endorse or

<sup>1</sup> Famously Descartes tries to find epistemic *terra firma* for inquiry independent of inquiry. But he doesn’t succeed. Thus “the Cartesian circle”, the problem of the foundationalist not being able to doubt everything and also engage in inquiry. In Descartes’ case, he isn’t able to doubt language, logic, the principle of clarity and distinctness and also to engage in the thought experiment that yields the *cogito*, the security of logic, and the principle of clarity and distinctness. The American pragmatists – specifically Peirce – recommends that we develop standards of “warranted assertability”, is a response to both Descartes’ problem, and to the additional or separate fact that science just never seems to yield certainty. One might say the shift from the *Theaetetus* account of knowledge in terms of justified true belief to a warranted assertability account – the distance between Plato and Peirce – was caused by science, by observing its history. Even the standards of epistemology itself are open to critique, revision, and expansion. Consider the evolution of the straight rule of induction as a tool of science. We start with the rule: if I observe that A & B co-occur to m/n, then infer that the next A will come with B to m/n. But we learn to insert rules in the antecedent such as so infer “if the sample is representative and large enough”. Eventually the straight rule yields, *a la Reichenbach* (1949), the canons of inductive logic, statistics and probability theory.

not endorse our own epistemic states, as the act of «placing [these states] within the logical space of reasons, of justifying and being able to justify what one says» (p. 76). What is this “logical space of reasons”? How, and where, is it, “reason”, situated in the four dimensional manifold of space-time, and how does it operate to produce claims that are of high quality and not merely of a certain quantity? Where in a world of causes is there space for reason?

The “task”, according to Sellars, is one of

[S]howing that categories pertaining to man as a *person* who finds himself confronted by standards (ethical, logical, etc.) which often conflict with his desires and impulses, and to which he may or may not conform, can be reconciled with the idea that man is what science says he is. (1963, p. 38)

In one guise, this problem is to reconcile the freedom in our lives suggested by the ability to decide what to believe with laws of nature that might seem to leave no room for decision or free action. In another guise, it is to account for the appearance of epistemic normativity within the causal, descriptive enterprise of the natural sciences. Failure to resolve this problem suggests that knowledge, as understood by epistemology, is mere appearance, and that our self-image as rational animals is an illusion.

The problem is made even more acute by the fact that the sciences themselves make knowledge claims. It is in this respect that the scientific image seems to be founded on, and thus incapable of replacing, what Sellars calls the “manifest image”, an ideology organized around a suite of concepts required for the operation of reason. The idea is that (a) the sciences are made intelligible as sciences by their ability to describe phenomena and evaluate claims about them, and (b) this ability has a parent in philosophical thinking’s broader freedom to call anything whatsoever into question. For the sciences to question and reject the ability to make up and change our minds about what to believe according to epistemic standards would be for them to undercut both their own intelligibility as a conglomeration of rational research programs, and, as a consequence, any epistemic authority those research programs might appear to have.<sup>2</sup>

<sup>2</sup> Hilary Putnam seems to have worries along these lines when he laments Quine’s call for naturalizing epistemology as “mental suicide” (1982, p. 20).

## 2. Ontological Commitment to the Kinds of Intentional Psychology

The problem of epistemic normativity is to make sense of the position from which we decide what to believe, and to identify the scope and limits of what empirical knowledge can tell us about the quality of those decisions. Sellars describes the problem as one of integrating two related logical frameworks that govern inquiry into the world and our place in it. These frameworks are his *manifest* and *scientific* images.

The manifest image is the system in which persons come to recognize themselves as persons, that is, as conscious, rational, free beings. It is organized around concepts used to both describe and evaluate various aspects of human life. The descriptive concepts are concepts like *perception, sensation, pleasure, pain, belief, desire, intention, goal, fear, hope, love, hate, choice, decision*. They all fall within the classificatory scheme of what is called intentional psychology as it is deployed within the manifest image to describe aspects of human life. The normative concepts are concepts like *good, evil, success, failure, justification, evidence, responsibility, duty, beauty, and wellbeing*. These are applied to assess and evaluate various qualities of the phenomena described by intentional psychology. These are second-order concepts in presupposing the legitimacy of those descriptions.

According to the manifest image, there are many kinds of cognitive entities that we are capable of considering, evaluating, and accepting or rejecting. At one end of the spectrum, there are phenomena like the contents of sensory experience, phenomena that Kant, for example, describes as what we receive, passively, in sensible intuition. At the other end of the spectrum, there are comprehensive ideologies: deeply held convictions, considered opinions, and whole systems of belief. To question these ideologies in deliberation is to consider the possibility that a whole way of understanding the world and our place in it is mistaken. Sellars reminds us in *Philosophy and the Scientific Image of Man* that the manifest image is an ideology like any other. Is there any basis for thinking that it is immune to doubt or rejection as inquiry advances?

If the inquiry is conducted by the sciences, there are indeed grounds for doubt about the survival of manifest image. Let us provisionally assume, with Sellars, that the sciences have final say about what is true. The scientific image departs from the manifest image at just the point where theoretical claims make reference to entities that have no basis in the concepts of familiar experience, among them the concepts of intentional psychology. If the sciences of human

behavior *can* proceed without deploying those concepts, the manifest image will have turned out to be an ultimately dispensable ideology. If the sciences *must* proceed without deploying those concepts, matters will be worse. The manifest image, as the home of both intentional psychology and epistemology, will have turned out to be mistaken, misleading, foolish, or fictional, not worth using if we wish to advance our conception of reality as it is. In hopeful moments Sellars conceives of the situation as one in which the ontology of intentional psychology (and of common sense, generally) is preserved as science provides “a needlepoint of detail” to that image; in other moods, he broaches the eliminativist possibility.<sup>3</sup>

### 3. Two Arguments Against the Possibility of Traditional Epistemology

The sciences motivate two lines of argument against the possibility of epistemology. The first of these threatens the legitimacy of intentional psychology and, in particular the concept of *belief*. Paul Feyerabend and Richard Rorty were the first to develop the idea, but the most straightforward version is an argument of Paul Churchland’s (Cf. Churchland, 1981).<sup>4</sup> Churchland takes seriously Sellars’s idea from *Empiricism and the Philosophy of Man* that intentional psychology is a theory in which thought is conceived by analogy to overt speech acts.<sup>5</sup> Churchland thinks it is a bad theory, on par with long-discredited Aristotelian physics and creationist biology, and thinks that the concepts of thought it deploys should and will be replaced by concepts from the neurosciences. He writes:

A look at the history of [intentional psychology] does little to allay such fears, once raised. The story is one of retreat, infertility, and decadence. The presumed domain of [intentional psychology] used to be much larger than it is. In primitive cultures, the behavior of most of the elements of nature were understood in intentional terms. The wind could know anger, the moon

<sup>3</sup> Sellars, in the sort of hopeful moment we have in mind, says this of philosophy in the Platonic tradition: «Let me elaborate on this theme by introducing another construct which I shall call – borrowing a term with a not unrelated meaning – the perennial philosophy of man-in-the-world. This construct, which is the “ideal type” around which philosophies in what might be called, in a suitably broad sense, the Platonic tradition cluster, is simply the manifest image endorsed as real, and its outline taken to be the large scale map of reality to which science brings a needle-point of detail and an elaborate technique of map-reading» (1963, p. 8).

<sup>4</sup> See Paul Feyerabend (1963) and Richard Rorty (1965,1970) for the earliest contemporary expressions of eliminative materialism.

<sup>5</sup> Sellars rehearses this view in *Philosophy and the Scientific Image*.

jealousy, the river generosity, the sea fury, and so forth. These were not metaphors. [...] To use Imre Lakatos' terms, [intentional psychology] is a stagnant or degenerating research program, and has been for millennia. (Churchland, 1981, p. 75)

The argument is that there may be no such things as beliefs, and, moreover, that there is *evidence* that there are no beliefs, as they are conceived by perennial philosophy. And if there are no beliefs, then there are no justified true beliefs, that is nothing to which concepts of justification and truth may be applied. There are no good witches and no wicked witches because there are no witches. One hope is that epistemology in the scientific image might recover something belief-like, something representational, that could play the role that beliefs conceived as discrete mental states play within the manifest image.

The second line of argument targets the justification relation. It is also very straightforward. It goes like so: What makes a science a science are its laws. Laws are counterfactual-supporting generalizations about causal relations. Therefore the only relations there are within the sciences are causal relations. On the assumption that the sciences have final say about what is true, causal relations are therefore the only relations there are; *sui generis* justificatory relations are an illusion.

As straightforward as this argument is expressed it is not as powerful as its proponent thinks. First, quantum physics questions the fundamentality of the causal relation in physics. Even if, as many say, causality is fundamental above the quantum level and even if the sciences above that level understand causal relations to be ontologically, or at least explanatorily, fundamental, there are nevertheless many non-causal relations that the sciences at least recognize, and, in some cases, find indispensable. There are spatial relations, expressed by claims like "Plato is to Aristotle's right", and "Michelangelo is painting upside-down again". There are temporal relations, expressed by claims like "Confucius was born before" Socrates. There are quantitative relations, expressed by claims like "Russell has more whiskey in his glass than Whitehead". There are mereological relations, expressed by claims like "Part of the painting is smudged". There are statistical relations, expressed by claims like "The majority of Leonardo's artwork is unfinished". This is not to mention all of mathematics: systems of relations that are (a) non-causal, and (b) theoretically indispensable to the natural and social sciences.

Despite all the non-causal relations the sciences either recognize or, with mathematics, presuppose, the thought remains that the sciences, as part of a descriptive enterprise, could never assess and prescribe epistemic standards. Jaegwon Kim (1988) gives a well-known version of this idea in objecting to W.V. Quine's suggestion in *Epistemology Naturalized* that «epistemology, or something like it, simply falls into place as a chapter of psychology and hence of natural science» (Quine, 1969). Kim objects to Quine's apparent recommendation to replace the study of a normative relation, justification, with the study of a causal relation, the one that holds between sensory input and verbal and written output. «[Quine] is asking us to set aside the entire framework of justification-centered epistemology. That is what is new in Quine's proposals. Quine is asking us to put in its place a purely descriptive, causal-nomological science of human cognition» (1988, p. 388). Kim's worry is that if the sciences have final say about what is true, and if all the sciences are authorized to speak about is the way things are, and why things are as they are, then they cannot (also) say anything about what we *ought* to believe, about the *norms* of belief, since *oughts* and *norms* are not the sorts of thing about which science, according to its job description, has, or is entitled to have, opinions about. Thus there can be no epistemology, traditionally conceived, within the scientific image.<sup>6</sup>

#### 4. Eddington's Tables

The surprising twist in the plot of scientific advancement is that the more we come to know about the world, the less we understand it. The less we understand it because the less our commonsense language seems to grasp things as they “really” are. This is why Sellars invokes Arthur Eddington's “two tables” when discussing “the clash of images” (1963, pp. 35–36) in *Philosophy and the Scientific Image of Man*. In the passage Sellars refers to, Eddington worries about the potential conflict between the world of everyday experience and the conclusions of scientific inquiry:

<sup>6</sup> One can dispute Kim's reading of Quine (Flanagan, 2006). The point of Quine's remark is that the (descriptive) project of rational reconstruction ought to be replaced by the (descriptive) project of neuropsychology, *not* that the *prescriptive* project of traditional epistemology ought to be replaced by the *descriptive* project of correlating sensory input with verbal and textual output. Nevertheless, it remains to be seen what the norms of naturalized epistemology come *to* if they are to come *from* the empirical sciences, given that the empirical sciences seem to be limited to descriptive claims.

I have settled down to the task of writing these lectures and have drawn up my chairs to my two tables. Two tables! Yes; there are duplicates of every object about me – two tables, two chairs, two pens. [...] One of them has been familiar to me from earliest years. It is a commonplace object of that environment which I call the world. How shall I describe it? It has extension; it is comparatively permanent; it is coloured; above all it is substantial. [...] My scientific table is mostly emptiness. Sparsely scattered in that emptiness are numerous electric charges rushing about with great speed; but their combined bulk amounts to less than a billionth of the bulk of the table itself. (1928, p. ix–x)<sup>7</sup>

We start in inquiry, eventually in the sciences, with what is most direct and familiar, only to find our initial understanding undercut by continual theoretical ramification and revision. This is bewildering enough when the subject is physics, but what happens when the sciences turn their attention to understanding and knowledge in their own right? It is one thing for the world at large to become more alien with the development of theory, it is another to form an alien self-conception, a picture in which we no longer understand ourselves to be who we thought we were. As the maturation of physical theory undercuts our familiar conception of the world, it stands to reason and is indeed happening (witness the revolutions in evolution, genetics, and now in neuroscience) that the scientific study of persons will not leave our own familiar self-conception intact, including the very concept of a person itself.

We can sum up so far: Preserving normative epistemology as more than a shadow of its former self demands two things: (1) finding room for the ontological categories of rational thought, or something close enough, within the scientific image; and (2) understanding how “believing” and its suite are more than just causally produced and causally productive informational states, but are actually knowledge-yielding. For our familiar self-conception to survive scientific advancement, the manifest and scientific images must share a common model of reason, what we have been calling critical thinking.

<sup>7</sup> For a similarly provocative comment on the “knowing” vs. “understanding” divide from a famous physicist, we have Richard Feynman saying in *The Character of Physical Law*:

«On the other hand, I think I can safely say that nobody understands quantum mechanics. So do not take the lecture too seriously, feeling that you really have to understand in terms of some model what I am going to describe, but just relax and enjoy it. I am going to tell you what nature behaves like. If you will simply admit that maybe she does behave like this, you will find her a delightful, entrancing thing. Do not keep saying to yourself, if you can possibly avoid it, “But how can it be like that?” because you will get “down the drain”, into a blind alley from which nobody has yet escaped. Nobody knows how it can be like that» (1967, p. 129).

Following Descartes we can also describe the capacity for reason as the faculty of judgment. Under any name, this capacity essentially involves the application of normative concepts to the psychological states postulated by intentional psychology. It is this that places them, as Sellars says, within the space of reasons. What does this capacity come to? How is it possible for material beings in a material world to execute this capacity (Flanagan, 2007)? Answering such questions involves at least two steps: first, we need to accurately describe the phenomenology of the space of reasons; second, we need to provide a naturalistic account that explains how the phenomenology can be realized by mammals like us.

## 5. The Phenomenology

John McDowell nicely describes the phenomenology of standing in the space of reasons in this passage:

For that kind of locution to fit, in the sense in which I intend it, the subject would need to be able to step back from the fact that it is inclined in a certain direction by the circumstance. It would need to be able to raise the question whether it *should* be so inclined, and conclude that it should. Acting on the inclination – supposing the verdict of the inquiry that is opened by this stepping back is positive – now takes on a dimension of freedom. (Unpublished, p. 7, Section 4)<sup>8</sup>

McDowell gives us the explanandum. What we need is a model of reason, immanent to the manifest image and portable to the sciences, and then refined perhaps by the sciences, but not overturned or eliminated by them, which explains how this phenomenology is possible. Methodologically, we proceed as Kant did, by starting with an adequate description of the phenomenology from which we can infer hidden cognitive processes. The methodological principle is that a system must be constructed such that it can perform the operations it actually performs. This is why it is so important to specify operational performance accurately. For the Kantian cognitive scientist, the good phenomenology functions to provide precise design specifications.<sup>9</sup>

<sup>8</sup> See Christine Korsgaard's *The Sources of Normativity* (2004) for a similar account of the phenomenology of the operation of reason.

<sup>9</sup> Flanagan identifies Kant as laying «both the substantive and methodological foundations for modern cognitive science» (1984, 1991). The transcendental deduction is a method of inferring underlying causes from overt phenomena. (Kant knew Newton's *Principia* very well, and it is likely that he was

## 6. How Epistemic Freedom is Possible

The next step is to provide a naturalistic account – one endorsed by the scientific image – that realistically models cognitive processes that might support the phenomenology of freedom to decide what to believe. Happily, we can help ourselves here to a solution that has been worked-out on the closely related problem of freedom in the domain of action.

Consider Daniel Dennett’s control-based approach to freedom of the will (1984). Dennett wants to explain a familiar aspect of human life – the apparent freedom we have to choose what to do – in light of information that constrains what that freedom could come to – laws of physics, for example, together with the now plausible, post-Darwin assumption that we are entirely material beings.

Dennett flags several major elements of the phenomenology of free choice as needing explanation: (1) Freedom of choice is experienced as coming in degrees: there are easy and difficult choices, no-brainers and effortful tasks, things I can do and things I can’t do (some of which I wish I could do); there are near goals and distant goals, and so on; (2) We can feel stuck or trapped when we have only bad options (or, as in the theory of cognitive dissonance, when we have multiple equally good, but incompatible, options; (3) We can be overwhelmed by having too many options to consider; (4) Normally, only some options are live options; and (5) Modal language about alternative possibilities, about whether and how things could be, and/or could have been otherwise, seems apt when speaking about choices-to-be-made or that-have-been-made.<sup>10</sup>

The first thing to note is the strong similarity between Dennett’s description of the phenomenology of free choice and McDowell’s account of how operating in the space of reasons can seem. The next thing to note is how

inspired by Newton’s own stated method of inferring hidden forces of nature from propositions about phenomena of motion. For more on Newtonian methodology, see Stein, 1990; Smith, 2007.

<sup>10</sup> Dennett analyses systems that possess a certain kind of natural freedom, i.e., freedom that is naturalistically possible, this way: “A system has a degree of freedom when there is an ensemble of possibilities of one kind or another, and which of these possibilities is actual at any time depends on whatever function or switch controls this degree of freedom. Switches (either on/off or multiple-choice) can be linked to each other in series, in parallel, and in arrays that combine both sorts of links. As arrays proliferate, forming larger switching networks, the degrees of freedom multiply dizzyingly, and the issues of control grow complex and non-linear. Any lineage equipped with such an array confronts a problem: What information *ought* to modulate passage through this array of forking paths in multi-dimensional space of possibilities? (2004, p. 162)”

we can learn from Dennett's move to answer the question: how is such freedom possible, if it is, in the natural world, the world described by science?

Dennett goes beyond the phenomenology in seeking to provide a model of freedom in terms of the operation of control over what Dretske (1983) calls the *flow of information*. What is control? Dennett writes:

The root idea of control, which has been elevated into a technically precise concept in cybernetics and automata theory, is (in ordinary terms) that *A controls B* if and only if the relation between *A* and *B* is such that *A* can *drive B* into whichever of *B*'s normal range of states *A wants B* to be in. (1984, p. 52)

The promise of Dennett's model of freedom and practical reason lies not only in its scientific viability, but in its being a model around which we can construct an account of theoretical reason, in which the flow of information is, or at least can be, controlled by consulting epistemic standards, or better, by processing through epistemic norms. Let us sketch out that account, starting with three aspects of McDowell's phenomenological story to focus on as design specifications to be met.

But first we must remind ourselves that critical reasoning, second-guessing, wondering, and changing beliefs may all be relatively rare cognitive processes. Many beliefs, perceptual ones, most familiarly, often just arrive, we go with them, and things work out. Usually when a perceptual experience gives us pause, makes us wonder, not all possibilities are entertained – we wonder whether that was a hawk or an owl, not whether that was some kind of bird, or some kind of kite, or some kind of extra-terrestrial, or some kind of hallucination. But sometimes the world or a text or a fellow inquirer or conversant gives us pause and makes us engage in wandering about the space of reasons. Then the familiar phenomenology ensues.

Here are some features of the phenomenology: First, the act of belief examination or, what is different, justification is experienced as some kind of wondering about or assessment of what one is, or was initially, inclined to believe, an evaluation of what is on the table as a candidate for epistemic endorsement. The cause of the wondering, the second-guessing and double-checking, is that something in the world has given one pause, has put the stops on automatic assent or dissent. Normally, when we are betwixt and between beliefs only some possible beliefs are live options. Considering a set of live options, plausible contenders for endorsement, is in fact, all it normally means to open an inquiry, or, alternatively, to have an open mind. The act of belief

examination or justification is really experienced as a process. The act is really an activity, experienced as an admixture of choate and inchoate, disciplined and undisciplined thoughts, resulting in a decision about what to believe for now.

## 7. Epistemology Naturalized

The free and critical operation of our rational capacities lie at the center of the manifest image endorsed by traditional philosophical thought (Cf. Flanagan, 2002). The topic of their formulation and legitimacy singularly define perennial philosophy as a series of footnotes to Plato, who was the first to treat them systematically. Consider the psychology of the *Republic*. In response to Adeimantus's request for a constructive defense of justice, Socrates suggests understanding the hidden activity of the soul by analogy to overt city life. He proceeds to identify reason with the rulers of a just *polis*, whose job is to weigh the preferences of the military and working classes against the best interest of the polis as a whole. The soul and the *polis* each operate rationally when alternative courses of action are considered for selection by the ruling class.

Plato does two very interesting things, here. He characterizes the soul in terms of parts that have functional configurations, and he identifies justice as the virtue realized by one configuration in particular. This configuration is the one in which the calculating part of the soul is free to assess the recommendations of the emotional and appetitive parts against its own concern for the whole. Plato's hypothesis is that *practical* reason functions to improving decisions about what to *do*. Our suggestion is that we think of *theoretical* reason, the ability to think critically about what to *believe*, by analogy to Plato's functional conception of practical reason, and if we can, is this a model of reason the sciences can accept?

In both metaethics and meta-epistemology, the issue is whether an ability, a sort of freedom to decide how to act and what to believe, respectively, remains available to us in light of information we have about how the world – including most relevantly, the mind – works. Some say knowledge about how the world works – discovery of the laws of physics and psychology – seems to crowd out deliberate action and thought. One can simply deny that we have any such freedom – what you find with hard determinism in ethics and the pure descriptivism Kim is worried about in epistemology – or make room for it by foisting a compatibilist view.

Our proposal is similar to John Dewey's – a compatibilist in ethics and epistemology – in analyzing the kind of freedom we care about in both ethics and epistemology as the ability to produce differential responses (beliefs and actions among them) to future circumstances based on feedback about past successes and failures.<sup>11</sup> To possess this sort of freedom is to be *response-able*, or *responsible* with an 'a'.<sup>12</sup> This is a modest conception of freedom, to be sure, but *if* it licenses an operational model of theoretical reason in the mold of Plato's account of practical reason, and if the sciences can accept a model of practical reason like Dennett's, this conception of freedom further licenses a model of reason that has potential staying-power inside the scientific image.<sup>13</sup>

The fact that the model of reason fundamental to the manifest image *may* be acceptable to the sciences, that it is not *inconsistent* with science, does not entail its usefulness (or indispensability) to them. Does the suite of concepts required for the operation of reason play a theoretical role within the scientific image? There are really two questions here. (1) Do the sciences apply an operational model of reason according to which we reflect upon, evaluate, and decide what to believe, and (2) do the sciences describe this operation as actually resulting in the acquisition of justified true beliefs?

The answer to (1) is yes, though by focusing on the use of intentional psychology in everyday experience – naive intentional psychology – critics have overlooked the extent to which the concepts of belief and desire (or some naturalistic descendants of them) are deployed within the scientific image. Recall Paul Churchland's comment that intentional psychology is a «stagnant or degenerating research program, and has been for millennia» (1981, p. 75). Churchland may be right that intentional psychology, or parts of it, did

<sup>11</sup> See Dewey (2002): «For morals has to do with acts still within our control, acts still to be performed. (p. 18) [...] The moral issue concerns the future. It is prospective. [...] The moral problem is that of modifying the factors which now influence future results» (p. 19). Flanagan discusses the sense in which reason is forward-looking in *The Science of the Mind* (1991, p. 51).

<sup>12</sup> See Flanagan's *The Really Hard Problem* (2007): «Dewey says the moral problem concerns the future. I treat you as an intelligent being, capable of self-control, if I call you on inappropriate or non-virtuous actions. If you are receptive and paying attention, my response gives you reason to behave better in the future. I call this *responsibility* to indicate that it incorporates the credible assumption that our characters, our hearts and minds, are plastic to some degree. Social communities are dynamic systems in which complex feedback mechanisms help us adjust our beliefs, desires, feelings, emotions, and behavior» (p. 35).

<sup>13</sup> Compare with Paul Churchland's (2012) view that reasoning is pattern and meta-pattern recognition. For a dissenting view, Alex Rosenberg (2011) continues to hold the line of eliminative materialism.

stagnate for millennia, but this stagnation abruptly ended with Pascal's inaugural treatment of probabilistic judgment in the 17th century, which continues to lead to advances in formal epistemology, broadly construed. Another research program that preserves – by bringing “needlepoint of detail” to – the concepts, categories, and processes that the manifest image endorses, is cognitive science and neuropsychology, a field that Churchland originally thought would eliminate the manifest image, but, in fact, has only sophisticated it.

Consider the field of formal epistemology. Pascal wanted to know how a person ought to choose from an array of uncertain options. He thought that two variables needed to be known to answer this question: the magnitude of an option's reward, and the probability of its being acquired. Pascal then identified the best option, the option that ought to be selected, as the product of those two variables. Importantly, he conceived of the objects of choice as objective quantities, in that the values of alternatives were not taken to be relative to a person's beliefs or desires. In fact, beliefs for Pascal only entered into the picture as potential *objects* of choice, which themselves could be evaluated as the product of their probability of being true and the reward a person would receive for holding them. Famously, Pascal claimed that choosing to believe in God is a better bet than not.

The major problem with Pascal's idea was that it did not adequately describe the decisions that people actually make. That is, even if events have objective probabilities of occurring, and even if those events have objective reward magnitudes, people systematically fail to choose options with the highest objective expected reward. In response to this *descriptive* failure, Daniel Bernoulli (1954) subjectivized Pascal's choice values by relativizing them to a person's expected (believed) utility (desirability). Bernoulli thus imported the notions of belief and desire from naive intentional psychology, quantified them as Pascal had quantified objective probabilities and reward values, and identified their product as the object of human decision making, what has become known as *subjectively expected utility*.

After Bernoulli, Pareto (1927/1971), von Neumann and Morgenstern (1944), and Samuelson (1937, 1938, 1983) successively developed and improved the concept of subjectively expected utility. Despite continual refinement, it remains a concept both (a) born and with roots in, naive belief-desire psychology, and (b) the central notion in economic theory, in social choice theory, and many parts of psychology.

More recently, Richard Jeffrey (1990) has proved that the decision matrices of decision theory, which were traditionally populated by goods and actions, can be interpreted to be populated by propositions or states of affairs, opening the door to a conception of decisions about what to believe as decisions between what we have called possible beliefs.<sup>14</sup> And in the last several years neuroscientists like Michael Shadlen and William Newsome (Cf. Shadlen and Newsome, 1996), and Paul Glimcher (Cf. Glimcher, 2010; Platt and Glimcher, 1999) have identified likely neural correlates of both subjectively expected utility values and the processes of choosing among options on the basis of those values.<sup>15</sup>

The point of this rush through the history of decision theory is that by developing a fully quantitative, scientifically respectable notion of rationality in terms of preference consistency, empirical theories of decision-making have been providing “needlepoint of detail” to the roughed-out design of intentional psychology handed off to the sciences by perennial philosophy. And so the proper reply to eliminative materialism is that intentional psychology survived into the scientific image as microeconomics, statistics and probability theory, game theory, and, lately, neuroeconomics, and neuropsychology, and that these disciplines remain secure within the scientific image under the umbrella of the decision sciences

What about (2), the question about whether the sciences describe the operation of reason as one that actually results in justified true beliefs, warranted beliefs, or something in their vicinity.? The answer to this question is complicated. The first complication is that there may be circumstances in which having true beliefs or acquiring information is a bad thing for you, if, that is, your primary interest is personal happiness. Under these circumstances, you are better off being under a *positive illusion* (Cf. Taylor and Brown, 1988; Flanagan, 2007; McKay and Dennett, 2009; Flanagan, 2009). To be better off under a positive illusions is to be better off maintaining an uncritical stance toward your own epistemic states, the risk in questioning them being the disruption of a way of life that (a) makes you happy, and (b) is unrecoverable once disrupted.<sup>16</sup>

<sup>14</sup> This is what attracts Donald Davidson to Jeffrey’s view in his own attempts to naturalize reason. See Davidson (1995) and (2004).

<sup>15</sup> Glimcher’s (2010) account of the history of economic thought and its recent integration with the neurosciences is especially illuminating.

<sup>16</sup> Flanagan (1991, 2007, 2009) argues that some positive illusions are not beliefs and thus are not illusions. These are better analyzed as hopes or wishes, possibly unrealistic hopes or wishes.

The second complication is that there may be circumstances in which more information is a good thing, but deliberating about what to believe is unlikely to deliver that information. Since the 1970s, we have become more familiar with the role in decision making of unconscious cognitive heuristics. Indeed many choices do seem to be better made without withdrawing from our circumstances to make them. Gerd Gigerenzer and his group, for instance have advanced a version of Herbert Simon's (1955) idea that the computational complexity of real-time decision making means that we are *in general* better off making choices about what to believe using "fast and frugal heuristics" (Cf. Gigerenzer and Todd, 1999).

There are nevertheless pitfalls to a systematic refusal to think about what to believe. Daniel Kahneman gives a favorite example from the work of his colleague Shane Frederick. Students at Princeton and the University of Michigan were given the following problem:

"A bat and a ball cost \$1.10 in total. The bat costs \$1 more than the ball. How much does the ball cost?" (2003, p. 1450)

The answer subjects are inclined to give, "10 cents", is wrong; and of course upon reflection we see quickly that the correct answer is "5 cents". But why the slip? Kahneman rehearses this study in his *Thinking, Fast and Slow* as an illustration of the function of our capacity for deliberation, what he calls System 2 cognition: «One of the main functions of System 2 is to monitor and control thoughts and actions "suggested" by System 1, allowing some to be expressed directly in behavior and suppressing or modifying others» (2011, p. 43) Kahneman argues that in absence of a prompt to double-check your own epistemic states, something signaling the need to scrutinize them, there is nothing to cause our brains to engage the metabolically costly routine of deliberating about what to believe.

The point can be made metaphorically as well. Why is it so hard to throw a Wiffle ball with any great velocity? The obvious answer is that a Wiffle ball is not heavy enough to throw very fast, and the obvious answer is correct. But it is not correct for obvious reasons. One may suppose that the lightness of the ball makes it more susceptible to drag forces, as a feather's lightness prevents it from falling to the earth with the acceleration of a bowling ball, and perhaps there is some truth to this idea. The real problem, however, does not have to do with drag forces. The problem is that Wiffle balls are not heavy enough to stimulate the full recruitment of the muscle fibers needed to generate the

power required to throw them with great velocity.<sup>17</sup> And just as some measure of resistance is required to recruit our most powerful skeletal muscles, so too some measure of resistance is required to recruit our most powerful mental muscles. Resistance, in the latter case, takes the form of feedback that leaves us not knowing what to believe.

Scientific advancement gives us feedback that leaves us not knowing what to believe about our own capacity for critical thought. This is far more significant than feedback that our first impressions are mistaken, or that more muscle fibers must be innervated to accelerate an object through space. But by characterizing this capacity operationally, in a way that runs back to Plato's understanding of critical thinking, we are coming to understand how a model of reason endorsed by perennial philosophy survives into the scientific image *as* a model of reason. If reason is the capacity to evaluate alternatives in light of the need for further information, and if reasoning about what to believe results in the acquisition of information, then critical thinking remains the sort of faculty that is good to have in situations where revising one's belief or acquiring more information in order to check one's beliefs is to one's advantage.

Our claim is that the sciences are equipped to make sense of reason, of our freedom to think critically and make up our minds about what to believe, if we extend Dennett's control-based account to cover decisions about what to believe. To do this we must (a) define the space of reasons as a matrix of possible beliefs;<sup>18</sup> (b) define critical thinking as the consultation of (potentially implicit) epistemic standards by which we evaluate the evidence for contender

<sup>17</sup> The example comes from Mark Rippetoe's *Starting Strength*. «This is due to several factors involving the physiology of skeletal muscle contraction, among them the fact that a very high velocity movement does not allow enough time for the nerves to recruit many of the components that contribute to muscle contraction. Like trying to throw a wiffleball [sic], a very light weight moving very fast does not provide enough resistance to push against effectively. A baseball is pretty good to throw, because it's just about the right weight to throw hard and fast. Power is at a maximum when throwing a 16 lb. shot, due to the combination of weight and velocity. But a great big rock would be too heavy to allow for the production of much power, because of the very slow velocity even a very strong man could produce. So the load must be optimum for power production» (2007, p. 176).

<sup>18</sup> Because we are concerned here with the psychological side of reasoning, we characterize "reasons" first internally, as psychologically real particulars, along Davidsonian lines. If I am wondering about the best way to get from Durham to Chapel Hill, the routes, whatever they are, that I am now considering are my reasons, the only ones I have or know about. But Google (formerly God in such thought experiments) has information about the very best route. It is an objective fact that I have reason to consult Google Maps. But if I don't know about Google Maps, this is an external reason, one that exists, but that I don't have (internally) as a reason.

beliefs; and (c) explain epistemology as the descriptive-genealogical explication and normative critique of epistemic standards we actually consult. Epistemic standards are the norms, rules, and principles governing the process of deciding what to believe, or, more directly in the province of epistemology itself, what epistemic norms and decision procedures to advocate and utilize. The best explanation of epistemic norms and decision procedures is that they are cognitive strategies developed over the course of history, through personal and social experiences, to guide and improve decisions about what to believe.

### 8. Epistemic Normativity

We've claimed that overcoming worries about the *bona fides* of epistemology in a world seen from the perspective of the scientific image requires two kinds of work. First, we need to provide a naturalistic model that plausibly describes and explains how the familiar phenomenology of critical reasoning is possible, indeed, why it is as it is. We've shown that a compatibilist model in epistemology, where freedom of thought and belief, freedom of reasoning, is modeled on a compatibilist solution that responds to parallel eliminativist worries about free action, can satisfy this demand. This is the descriptive-genealogical part of the project. It paves the way to address the second part of the problem, the problem of normativity. This is the task of showing epistemology to be capable of discovering, expanding, critiquing, and endorsing the decision procedures that the sciences both utilize and describe and explain.

The epistemic normativity problem seems especially difficult. Why? Because neither the ability to *describe* aspects of human life *as* thought, nor the ability to describe thought as *deliberate* or *critical*, seems to license the ability to *evaluate* the standards employed in deliberation and critical thinking. This is Kim's worry about what the program of naturalized epistemology comes to, an apparent recapitulation of the ethical *is-ought* problem, which Sellars notices and mentions *as* a recapitulation of the ethical *is-ought* problem at the end of *Philosophy and the Scientific Image of Man* (1963, p. 39). The question, then, is how to get claims about the epistemic standards a person ought to consult from the sciences (or from philosophy conceived naturalistically, as continuous with science or as beholding to the scientific image) if all the sciences can do is set out descriptive-genealogical-explanatory claims. This worry does not depend on, and thus does not require therapy to overcome, the false belief that

the causal relation is the only legitimate relation. It depends, at this point in the dialectic, on the beliefs that science only traffics in description, genealogy, and explanation and that these do not yield oughts. Both beliefs are false.<sup>19</sup>

To see this, consider Quine's own response to the objection that in *Epistemology Naturalized* he had set out to abandon epistemic normativity:

Naturalization of epistemology does not jettison the normative and settle for the indiscriminate description of ongoing procedures. For me, normative epistemology is a branch of engineering. It is the technology of truth-seeking, or, in more cautiously epistemological term, prediction. Like any technology, it makes free use of whatever scientific findings may suit its purpose. It draws upon mathematics in computing standard deviation and probable error and in scouting the gambler's fallacy. It draws upon experimental psychology in exposing perceptual illusions, and upon cognitive psychology in scouting wishful thinking. It draws upon neurology and physics, in a general way, in discounting testimony from occult or parapsychological sources. There is no question here of ultimate value, as in morals; it is a matter of efficacy for an ulterior end, truth or prediction. The normative here, as elsewhere in engineering, becomes descriptive when the terminal parameter is expressed. (1986, pp. 664–665)

The key idea is this: If we stipulate that “truth” is the end, “the terminal parameter”, either given by our platonic, with a small ‘p’, nature, or by the set of social practices we call “science”, then epistemology is the practice of discovering, developing, critiquing, and endorsing the norms that produce it, “the terminal parameter”, truth.

Normative epistemology survives into the scientific image as a body of informed norms that prescribe epistemic decision procedures that reliably produce truth (or, at least, have done so thus far).<sup>20</sup> Our decisions about what to believe are good decisions if they meet criteria for how we ought to make up

<sup>19</sup> The key is to deny the antecedent of the conditional in the claim that: “if all the sciences can do is set out descriptive-genealogical-explanatory claims”, then no normativity can emerge. Science can extract norms, not demonstratively or deductively, but abductively, for its own practices that yield knowledge. When the scientific community does this, it goes 2<sup>nd</sup> order and does what we call “meta-science”. When statisticians, decision scientists, logicians, and philosophers engage in the same activity it is broadly the science of reasoning or, for simplicity, epistemology. The norms are both used to make first order claims and extracted from sets of such claims. Vicious circularity is avoided.

<sup>20</sup> See Flanagan (1982, 1988, 2006), where naturalized epistemology is set out as «the enterprise of sorting reliable techniques of knowledge acquisition from unreliable ones in theoretical domains» (p. 541). See also Goldman (1986), Nozick (1994), Lauden (1990), Kitcher (1992), Rosenberg (1990), and Kornblith (1993).

and change our minds, and the criteria for how we ought to make up and change our minds depend on the canons of epistemology conceived along the lines just articulated. Reichenbach's (1949) pragmatic justification of induction (see footnote 1), like the research program initiated by Pascal, is an example of epistemology in just this spirit. The norms of knowing have been refined in epistemology proper, as well as in the epistemic sciences that include statistics and probability theory, game theory, neuroeconomics, and everything else under the umbrella of the decision sciences. One ought not think that naturalized epistemology is in its early stages. It is not. Indeed, the best way to allay fears about its *prospects* is to come to appreciate its *history* as involving a large and precise body of norms that yield knowledge in the sciences, and in ordinary life.<sup>21</sup>

## 9. Conclusion: Language and the Space of Reasons

So, the scientific image endorses and sophisticates, rather than undermines or eliminates, the manifest image of persons as creatures who operate inside the space of reasons. What does the scientific image say about the relationship between epistemic agency and language? Are humans the only epistemic agents, the only creatures who traffic in reasons? What about nonhuman animals, infrahumans (e.g., infants), and extrahumans – the World Wide Web, the iCloud? From an empirical point of view, what degree of epistemic agency is language necessary for, exactly?

<sup>21</sup> Here's a thought for unification of ethics and epistemology that might be attractive to the naturalist – defense would require another paper. Insofar as epistemology and ethics involve imperatives, they are all hypothetical imperatives. The "insofar" is important. The ends of epistemology and ethics, might be categorical ends or goods, "truth" for epistemology, "good" for ethics (Flanagan, 2007). The idea that epistemology survives into the scientific image on the back of an instrumental conception of epistemic normativity might be thought to echo Philippa Foot's thesis that morality is a system of hypothetical imperatives (1972). Early Foot thought that moral norms are like other norms in being means to ends. She came later to explain the apparent necessity of moral norms as rooted in the *de facto* ubiquity and necessity (i.e., mandatory nature) of some ends (see *Natural Goodness* 2001). If we put both strands together, something Foot did not think could be done smoothly, ethics is naturalized by conceiving of it as an inquiry into the sorts of things we ought to do if we want to live a rewarding life as a *eudaimon*, where the exact causes and constituents of *eudaimonia* are open to empirical inquiry. Likewise, epistemology naturalized can be understood as the set of instrumental norms that serve the unconditional end of knowing. All the norms are hypothetical, where the antecedent sets the task as truth-seeking. The end of truth-seeking, the command, if you will, to seek truth is itself not conditional, hypothetical, or instrumental. The truth is good in itself.

There is a line of thought in Sellars's work, one taken up by Brandom (cf. 1995), and McDowell (cf. 1994; Unpublished), according to which language is necessary for possession of the suite of concepts required for the operation of reason and critical thought. Sellars, Brandom, and McDowell all acknowledge a continuity between non-linguistic animals and human beings but deny that non-linguistic animals can exhibit legitimately critical thinking. They cannot exhibit legitimately critical thinking because, *ex hypothesi*, they cannot, without language, bring into view for evaluation the very standards they use to evaluate their epistemic states and decide what to believe. The claim, then, would be that language, while not necessary for knowledge of the world in some attenuated, "as-if" sense of knowledge, is necessary for the evaluation and adjustment of the norms according to which such knowledge, or proto-knowledge, is acquired. But this claim is more than the claim that language is necessary for being a participant in the space of reasons; it is a claim about engagement in the special critical reasoning practice of epistemology.

Let us distinguish three phenomena: critical thinking or belief acquisition and revision; epistemology, or critical thinking about, or belief formation and revision with respect to, the standards by which beliefs are acquired and revised; and meta-epistemology, or critical thinking about, or belief formation and revision with respect to, the scope and limits of epistemology. Meta-epistemology involves scrutiny of deeply held convictions, considered opinions, and whole systems of belief: what we have previously identified as comprehensive ideologies.

Critical thinking, as we have described it, is a widespread phenomenon. Furthermore, only critical thinking is necessary for being an epistemic agent. Non-human primates and infants, when provided with environmental feedback in the form of cues and prompts for greater scrutiny, what we earlier called "resistance", are capable of checking, second-guessing, even double-checking their environments to acquire more information relevant to a surmise and in order to reach secure conviction. Dogs sometimes come to make sure a toy has been thrown before darting off for it. Any animal that lives in an environment with variable opportunities for food and sex must learn from experience about these opportunities in order to increase its likelihood of surviving, and thus of passing its genes to successive generations. Learning by checking surmises, by bringing epistemic states into view for endorsement, is sufficient to place such surmises, expectations, and other epistemic states within the space of reasons.

Non-linguistic animals cannot do epistemology as it is practiced in philosophy departments, but can they do anything beyond merely double-checking and second-guessing when confronted with reasons for greater scrutiny? We offer that any animal that has the ability to improve its learning strategies through experience, for example by becoming increasingly wary of small-sample sizes while foraging and thus behaving as though has adopted a new and improved version of Reichenbach's straight rule that has adjusted for sample size is doing something more complex than ordinary critical thinking. Darwinian gradualism encourages us to call such normative refinement among smart non-human animals, "proto-epistemology".

What language is undoubtedly necessary for is meta-epistemology, the central problem of which Roy Wood Sellars identifies as one of connecting knowledge up with the world, which the sciences study. This problem becomes urgent in *Philosophy and the Scientific Image*, where Wilfrid Sellars brings into relief the possibility that the very categories upon which the practice of epistemology depends may not themselves have a place within the scientific image. We have argued that even as the categories of intentional psychology and normative discourse are quantified and formally refined, they are not eliminated. The categories of true belief, justification, and warranted assertability, categories required to make sense of the sciences as sciences, have in fact survived into a multifaceted program of naturalized epistemology that is well underway.

#### REFERENCES

- Bernoulli, D. (1954). Exposition of a new theory on the measurement of risk. *Econometrica: Journal of the Econometric Society*, 22(1), 23–36.
- Brandom, R. (1995). Knowledge and the social articulation of the space of reasons. *Philosophy and Phenomenological Research*, 55(4), 895–908.
- Churchland, P.M. (1981). Eliminative materialism and the propositional attitudes. *The Journal of Philosophy*, 78(2), 67–90.
- Churchland, P.M. (2012). *Plato's Camera: How the Physical Brain Captures a Landscape of Abstract Universals*. Cambridge, MA: MIT Press.

- Davidson, D. (1995). Could there be a science of rationality?. *International Journal of Philosophical Studies*, 3(1), 1–16.
- Davidson, D. (2004). A unified theory of thought, meaning, and action. In D. Davidson, *Problems of Rationality*. Oxford: Oxford University Press.
- Dennett, D.C. (1984). *Elbow room: The varieties of free will worth wanting*. Cambridge, MA: MIT Press.
- Dennett, D.C. (1996). *Darwin's dangerous idea: Evolution and the meanings of life*. New York: Simon and Schuster.
- Dennett, D.C. (1997). *Kinds of minds: Toward an understanding of consciousness*. New York: Basic Books.
- Dennett, D.C. (2004). *Freedom evolves*. ePenguin.
- Dewey, J. (1921/2002). *Human nature and conduct*. New York: Dover Publications.
- Dretske, F. (1983). *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.
- Eddington, A.S. (1928). *The nature of the physical world*. New York: MacMillan Company.
- Feyerabend, P. (1963). Materialism and the mind-body problem. *The Review of Metaphysics*, 17(1), 49–66.
- Feynman, R.P. (1967). *The character of physical law*, Cambridge, MA: MIT press.
- Flanagan, O. (1988). Pragmatism, ethics, and correspondence truth: response to gibson and quine. *Ethics*, 98(3), 541–549.
- Flanagan, O. (2006). Varieties of naturalism. In P. Clayton (Ed.), *Oxford handbook of religion and science*. New York: Oxford University Press, 430–452.
- Flanagan, O.J. (1982). Quinean ethics. *Ethics*, 93(1), 56–74.
- Flanagan, O.J. (1984–1991). *The science of the mind*. Cambridge, MA: MIT Press.

- Flanagan, O.J. (2002). *The problem of the soul: Two visions of mind and how to reconcile them*. New York: Basic Books.
- Flanagan, O. J. (2007). *The really hard problem: Meaning in a material world*. Cambridge, MA: MIT Press.
- Flanagan, O.J. (2009). “Can do” attitudes: Some positive illusions are not misbeliefs. *Behavioral and Brain Sciences*, 32(06), 519–520.
- Foot, P. (1972). Morality as a system of hypothetical imperatives. *The Philosophical Review*, 81(3), 305–316.
- Foot, P. (2001). *Natural Goodness*. Oxford: Oxford University Press.
- Gigerenzer, G., & Todd, P.M. (1999). Fast and frugal heuristics: the adaptive toolbox. In G. Gigerenzer, P.M. Todd & the ABC Research Group, *Simple heuristics that make us smart*. New York: Oxford University Press, 3–34.
- Glimcher, P.W. (2010). *Foundations of neuroeconomic analysis*. Oxford: Oxford Univ Press.
- Goldman, A.I. (1986). *Epistemology and cognition*. Cambridge, MA: Harvard Univ Press.
- Jeffrey, R.C. (1990). *The logic of decision*. Chicago: University of Chicago Press.
- Kahneman, D. (2003). Maps of bounded rationality: Psychology for behavioral economics. *American Economic Review*, 93(5), 1449–1475.
- Kahneman, D. (2011). *Thinking, fast and slow*. New York: Farrar Straus & Giroux.
- Kim, J. (1988). What is “naturalized epistemology”? *Philosophical perspectives*, 2, 381–405.
- Kitcher, P. (1992). The naturalists return. *The Philosophical Review*, 101(1), 53–114.
- Kornblith, H. (1993). Epistemic normativity. *Synthese*, 94(3), 357–376.
- Korsgaard, C.M. (2004). *The sources of normativity*. Cambridge: Cambridge University Press.

- Laudan, L. (1990). Normative naturalism. *Philosophy of Science*, 57, 44–59.
- McDowell, J. (1994) *Mind and World*. Harvard University Press
- McDowell, J. (Unpublished). *Sellars and the space of reasons..* Retrieved from <http://www.pitt.edu/~brandom/mc-core/downloads/McD%20Cape%20Town%20talk--Sellars%20EPM.doc>
- McKay, R.T., & Dennett, D.C. (2009). The evolution of misbelief. *Behavioral and Brain Sciences*, 32(6), 493–561.
- Nozick, R. (1994). *The nature of rationality*. Princeton: Princeton University Press.
- Pareto, V. (1927/1971). *Manual of political economy*. Tr. by A.S. Schwier. New York: Augustus M. Kelly.
- Platt, M. L., & Glimcher, P.W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, 400(6741), 233–238.
- Putnam, H. (1982). Why reason can't be naturalized. *Synthese*, 52(1), 3–23.
- Quine, W.V.O. (1960). *Word and Object*. Cambridge, MA: MIT Press.
- Quine, W.V.O. (1969). Epistemology Naturalized. In W.V.O. Quine, *Ontological Relativity and Other Essays*. New York: Columbia University Press, 69–90.
- Quine, W.V.O. (1986). Reply to Morton White. In L.E. Hahn & P.A. Schilpp (Eds.), *The Philosophy of WV Quine*. The Library of Living Philosophers, vol. XVIII. La Salle, IL: Open Court Publishing Company, 663–665.
- Reichenbach, H. (1949). *The Theory of Probability*. Berkeley and Los Angeles: University of California Press.
- Rorty, R. (1965). Mind-body identity, privacy, and categories. *The Review of Metaphysics*, 19, 24–54.
- Rorty, R. (1970). In defense of eliminative materialism. *The review of metaphysics*, 24, 112–121.
- Rosenberg, A. (1990). Normative naturalism and the role of philosophy. *Philosophy of Science*, 57, 34–43.

- Rosenberg, A. (2011). *The Atheist's Guide to Reality: Enjoying Life Without Illusions*. New York and London: WW Norton & Company.
- Samuelson, P.A. (1937). A note on measurement of utility. *The Review of Economic Studies*, 4(2), 155–161.
- Samuelson, P.A. (1938). A note on the pure theory of consumer's behaviour. *Economica*, 5(17), 61–71.
- Samuelson, P.A. (1983). *Foundations of economic analysis*. Cambridge, MA: Harvard University.
- Sellars, R.W. (1922). *Evolutionary naturalism*. Chicago and London: Open Court Publ. Co.
- Sellars, W. (1963). Philosophy and the scientific image of man. In W. Sellars, *Science, perception, and reality*. London: Routledge & Kegan Paul, 1–40.
- Sellars, W., Rorty, R., & Brandom, R. (1997). *Empiricism and the Philosophy of Mind*. Cambridge, MA: Harvard University Press.
- Shadlen, M.N., & Newsome, W.T. (1996). Motion perception: seeing and deciding. *Proceedings of the National Academy of Sciences*, 93(2), 628–633.
- Simon, H.A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1), 99–118.
- Smith, G. (2007). Newton's philosophiae naturalis principia mathematica. An entry in the on-line *Stanford Encyclopedia of Philosophy* <<http://plato.stanford.edu>>
- Stein, H. (1990). "From the phenomena of motions to the forces of nature": Hypothesis or deduction? *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 209–222. Available from <<http://www.jstor.org/stable/193069>>
- Taylor, S., & Brown, J. (1988). Illusion and well-being: A social psychological perspective on mental health. *Psychological Bulletin*, 103(2), 193–210.
- Von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton: Princeton University Press.