Social Epistemology as the Science of Cognitive Management: Releasing the Hidden Potential in the History of Philosophy

STEVE FULLER

1. Social Epistemology as Cognitive Economy

Twenty-five years ago, in the first edition of *Social Epistemology* (Fuller, 1988), I began as follows:

The fundamental question of the field of study I call social epistemology is: *How should the pursuit of knowledge be organized, given that under normal circumstances knowledge is pursued by many human beings, each working on a more or less well-defined body of knowledge and each equipped with roughly the same imperfect cognitive capacities, albeit with varying degrees of access to one another's activities?*

It is clear from this form of words that I envisaged social epistemology as a kind of 'cognitive management'. The titles of some of my later books, such as *The Governance of Science* (Fuller, 2000a) and *Knowledge Ma-*

nagement Foundations (Fuller, 2002), suggest as much. However, the spirit of this enterprise differs from that of what is normally called 'cognitive science', which as Jerry Fodor (1981) shrewdly (yet approvingly) observed, assumes a Cartesian starting point that would have us understand the mind in its own terms before trying to figure out its relationship to the non-mental world. Thus, 'artificial intelligence' has been more concerned with specifying the conditions that would qualify an entity as 'intelligent' than with whether such an entity must be an animal operating in physical environment or can be an simply avatar in cyberspace.

In contrast, my version of social epistemology considers, so to speak, the 'formal' and 'material' elements of cognition at the same time and, in that respect, it is closer to *economics* in its conception. By that I mean that whatever cognitive goals we may wish to pursue, we need to consider the costs, how those costs would be borne and, as a consequence, whether the goals are really worth those costs. While this economic specification gives social epistemology a concreteness that has been often lacking in contemporary theories of knowledge, it by no means involves a downsizing of our epistemic ambitions. It is simply a call for those engaged in 'knowledge policy' (a term coined in Fuller [1988]) to provide an open balance sheet that reveals the costs and benefits behind particular strategies of cognitive re-organization. We may indeed be willing to suffer radical changes to our lifestyles and work habits, if we think a particular set of goals are worth pursuing. But wherever there is a gap, the social epistemologist has her work cut out.

In the back of my mind when I wrote those opening words in 1988 was Adam Smith's argument for the rationalisation of the division of labour in the economy as a means to increasing society's overall wealth. Smith observed that individuals doing everything for themselves were less efficient than each person specializing in what they do best and then engaging in exchange with others to obtain what they need. My point here is not to endorse any specific policies inspired by Smith but to acknowledge that he thought about the matter the right way in the following two senses:

(1) people are capable of changing even their fundamental habits if provided with sufficient reason (or 'incentive');
(2) people are a source of untapped potential that may be released by altering ('liberalising') the conditions under which they are allowed to express themselves.

Many things are implied here, perhaps most importantly the plasticity of human beings and hence the openness to social experimentation. Human history has only revealed a fraction of what

we are capable of. This is a faith that united both capitalism and socialism in the modern era – and one that my version of social epistemology carries forward.

Perhaps in these 'times of austerity', the drive to 'economize' is understood as a counsel to 'do more with less' in a way that presupposes that we have fewer resources than we first thought. On the contrary, when Smith and the original political economists in Britain and France promoted 'economizing' in the 18th century, they had in mind working more efficiently so as to conserve effort so that more can be done. This is the context in which greater productivity was seen as a natural consequence of the rational organization of human activity (Rothschild, 2001). We are held back not by the finitude of matter but the finitude of our minds to manage matter. The benchmark for this entire line of thought was the Augustinian doctrine of *creatio ex nihilo*: The ultimate rationality of divine creation is that God creates everything out of nothing – that is, no effort is wasted whatsoever. And if we are created 'in the image and likeness' of this deity, which Augustine emphasized as a lesson of Genesis, then we are tasked with achieving this divine level of performance.

It is also worth distinguishing my version of cognitive management differs from the appeal to economics made by 'analytic social epistemologists', such as Alvin Goldman and Philip Kitcher, who for the past twenty years have gravitated to aspects of economics that play to their default methodological individualism, whereby knowledge is sought or possessed in the first instance by individuals and then aggregated into 'social knowledge' in a literal 'marketplace of ideas' (Fuller, 1996). Thus, analytic social epistemologists have fancied microeconomic models that propose the optimal flow of information, division of cognitive labour, etc. In contrast, my own sense of cognitive management concerns the *macro*economics of knowledge, which is concerned with the overall efficiency of the epistemic enterprise, what Nicholas Rescher (1978), with a nod to the US pragmatist philosopher Charles Sanders Peirce, properly called 'cognitive economy'.

The idea of 'cognitive economy' was a product of the so-called 'marginalist revolution' in the final quarter of the 19th century, when the study of political economy came to acquire the shape of the discipline that today we call 'economics' (Proctor, 1991: chap. 13). Peirce extended what had been the key conceptual innovation of that revolution: namely, the principle of diminishing marginal utility. Applied to knowledge production, this principle implies that the indefinite pursuit of a particular intellectual trajectory is justifiable not as an end in itself but only on a benefit-to-cost basis. Our best epistemic enterprises provide the most cognitive

benefit at the lowest cost. This principle was explicitly proposed for science policy by the 'finalization' movement associated with Jürgen Habermas when he directed a Max Planck Institute dedicated to the' techno-scientific life-world' in the 1970s (Schaefer, 1983). Their idea was that puzzle solving in Kuhnian normal science eventually suffers from diminishing marginal returns on further investment. Thus, rather than following the Kuhnian strategy of running paradigms into the ground by deploying enormous effort to make relatively little technical progress (which finally forces even the most dogmatic scientist to realize that a radical change in perspective is needed), the finalizationists beyond a certain point would shifted resources to fields with better epistemic yields or these mature fields would be drawn together to solve standing social problems – such as cancer or environmental degradation – that escape the expertise of any particular discipline.

However, ideas surrounding cognitive economy may be deployed in other ways, such as a principle for the critical evaluation of existing knowledge systems. Across the range of national and corporate research systems, the rate of return on investment varies significantly. For example, the US may by far produce the most science, but the UK is much more productive relative to resource allocation. A comparable point may be made about educational systems. Harvard and Oxford may produce the most impressive roster of graduates but they also have the most impressive intake of students. The added value, cognitively speaking, of attending these institutions is probably much less than universities operating with much fewer resources that nevertheless produce distinguished graduates out of students of humbler origins. Worth stressing is that the main value associated with cognitive economy in keeping with the Augustinian point about *creatio ex nihilo* is best measured in terms of the opportunity costs that can be minimised or avoided, as efficiency savings make more resources available for other projects. The underlying intuition is that one acts now so as to maximise the degree of freedom that is later at one's disposal. I have been toying with this idea for a while, originally as 'epistemic fungibility' (Fuller, 2000a: chap. 8).

2. Two Kinds of Cognitive Economy for Social Epistemology

To understand the dynamic of the history of epistemology as a species of cognitive economy, we need to start by distinguishing *demand*- and *supply*-side epistemic attitudes. Demand-siders

proportion their belief to the need served by the belief. In other words, the more necessary the belief is to one's sense of self, the more it will be actively pursued. In contrast, supply-siders believe in proportion to the available evidence for the belief, even if that leads to a more diminished sense of self. Demand-siders characteristically hold that knowing is not complete without doing (i.e. generating the knowledge products that satisfy our cognitive needs), whereas supply-siders typically put in less effort in the cognitive process and expect less in return (i.e. conserving what is already known and ensuring that it does not deteriorate or become contaminated). As a first approximation, the demand-sider might be regarded as holding an 'industrial' model of cognitive economy that is focused on increased productivity, whereas the supply-sider holds a more 'agricultural' model that is more concerned with a steady yield in balance with the environment.

To make this distinction still more vivid, consider the demand-sider as someone who treats his ideas as opportunities to formulate hypotheses that then lead him to conduct experiments to discover something about the world that he had not previously known, which then forces him to redefine his objectives. Such a person is clearly in the business of self-transcendence. Whether his experiments have turned out to be true or false, he has acquired a power that he previously lacked. The only question is whether he has budgeted properly to reap the full benefits of that potential. This 'budgeting' should be understood in both cognitive and material terms. In particular, the demand-sider needs to be flexibly minded to see the intellectual possibilities that are opened up by being forced to give up old epistemic assumptions as a result of an unexpected research outcome. To the supply-sider, this requires the remarkable capacity to remain mentally invested in an array of possible futures, including ones that go against most of one's own previous cognitive and material investments. Only a deity could be so capable such equanimity in the face of what are bound to be many thwarted expectations. In humans such an attitude can easily look like that of Dr Pangloss, Voltaire's satirical portrayal of Leibniz in *Candide*. Worse still perhaps, the supply-sider might wonder whether the demand-sider has not succumbed to what social psychologists call 'adaptive preference formation', specifically the kind that Jon Elster (1983) dubbed 'sweet lemons'. This is the inverse of 'sour grapes', whereby one becomes incapable of facing failure on its own terms, always seeing the silver lining in every cloud. In the course of this self-delusion, so the supply-sider worries, the demand-sider detaches himself from any sense of security and becomes reckless with his own life – and perhaps the lives of others.

At this point, it is worth remarking that what in a comic frame might appear panglossian, in a tragic frame might come to be seen in Nietzsche's Zarathustrian terms: 'What doesn't kill me makes me stronger'. (Stanley Kubrick's *Dr Strangelove* may be seen as someone whose identity shuttles between these two frames.) One contemporary context for understanding these two attitudes is former market trader's Nicholas Taleb's (2012) distinction between 'fragile' and 'antifragile' approaches to life, which correspond, respectively, to the world-views of the supply- and demand-side epistemologists. Taleb generalises the lesson that he first taught concerning 'black swans', namely, those highly improbable events that when they happen end up producing a step change in the course of history (Taleb, 2007). His starting point is a dismissal of those who claim in retrospect that they nearly predicted such events and think they 'learn' by improving their capacity to predict 'similar events' in the future. Such people, who constitute an unhealthy proportion of pundits in the financial sector (but also a large part of the social science community), are captive to a hindsight illusion that leads them to confuse explanation with prediction. The lesson they should learn is that prediction of extreme events is always a mug's game. Rather, what matters is coming out stronger regardless of how one's future predictions turn out.

In Taleb's presentation, antifragility belongs to a tripartite distinction in world-views, roughly defined in terms of how one deals with error or unwanted situations more generally. The 'fragile' agent is one who needs to control the environment in order to maintain its normal condition. A slight shift in the environment can result in devastating consequences. In supply-side epistemology terms, this is the problem of scepticism. In contrast, the 'robust' agent maintains its normal condition in response to changes in the environment. But an 'antifragile' agent always maintains or improves its current condition as the environment changes, without any preordained sense of normality. A sense of the difference between a 'robust' and an 'antifragile' agent is captured by, on the one hand, a gambler who is simply concerned with always being able to return to the casino no matter how his bets turn out and, on the other, a gambler who always bets so that his losses can never outpace his wins, which generally means placing a somewhat larger than expected bet on improbable events and a somewhat smaller than expected bet on probable ones. The robust gambler does it as a hobby, the antifragile one does it to make a living.

The key to the antifragile mentality is what Taleb calls 'optionality', namely, the use of degrees of freedom as a proxy for

knowledge. In other words, if you do not know what will happen, make sure you have most options covered. In gambling circles, it is called 'spread betting', and there is an art to exactly how much one should underestimate continuity and overestimate rupture with the past in order to profit significantly in the long term. Interestingly, some computer scientists hypothesize that intelligence dawns in physical systems that conserve their potential, neither by responding similarly to all contingencies nor by trying to limit the contingencies to which they are exposed. Rather, intelligence emerges from keeping as many options open as possible so that the agent flourishes regardless of the contingency encountered (Wissner-Gross et al., 2013). In practice, this implies a regular process of sorting the wheat and chaff in one's cognitive horizons – that is, distinguishing the features that need to be preserved in any possible future from those that may be abandoned once they appear to be a liability, thereby resulting in a sense of 'sustainable error'.

In any case, this process is psychologically much more difficult than it seems for two reasons, one obvious and the other subtle. Obviously, as the supply-side epistemologist would stress, much of our sense of reality's stability rests on the future continuing the past being a 'sure bet'. Why then waste time and money on outliers? Nevertheless, Taleb counsels that it is better to run slightly behind the pack most of the time by devoting a small but significant portion of your resources to outliers, because when one of them hits, the rewards will more than make up for the lower return that you had been receiving to date. This raises a subtler psychological difficulty with antifragility: Once you decide that your bets require redistribution – say, in light of failed outcomes – how do you preserve the information that you learned from your failed bets in your next portfolio of investments? Rarely is the matter as straightforward as simply shifting out of the failed bets to the ones that did best, since the latter may be only temporarily protected from the same fundamental problems that led to your other bets to fail. In other words, every failure provides an opportunity for a fundamental re-think about all your bets, including the successful ones. This is how 'learning', properly speaking, is distinguished from mere 'surviving' over time. In that sense, you really never reduce uncertainty but you learn to game it better.

Taleb's (2012) main piece of advice here is that one's epistemological insight is sharpened by having 'skin in the game', to use the gangster argot for having a material investment in the outcomes. Scornful of academic and other professional pundits, who are paid to issue predictions but are not seriously judged on their accuracy, Taleb dubs them the 'fragilista' because they are

insulated from the environments to which they speak. Thus, they have the luxury of behaving either like 'foxes' or 'hedgehogs' in the political psychologist Philip Tetlock's (2005) sense: that is, they can simply mimic the trends or stick with the same position until it is borne out by events. They have no incentive to think more deeply about the nature of the reality that they are trying to predict.

3. The History of Epistemology as a Struggle over Cognitive Economy

Immanuel Kant originally glimpsed the demand- and supply-side epistemic attitudes toward the management of knowledge production at the end of modern epistemology's cornerstone work, *Critique of Pure Reason* (1781). In that work, demand- and supply- side epistemology are famously canonised as representing two traditions with deep historical roots. They continued to be enshrined in the curriculum as the foundations of what is still called 'modern philosophy'. The demand- and supply-side attitudes are known, respectively, as *rationalism* and *empiricism*. Kant suggested that this distinction had been played out across the entire history of philosophy, moving roughly from one of general metaphysics to a more narrowly epistemological horizon, as the distinctness of 'the human' itself came more clearly into view. In Figure 1, I have elaborated the historical trajectory that Kant leaves implicit, by tracing the path of these parallel legacies from their classical expression in Plato and Aristotle through the alternative Hellenistic life-philosophies of the Stoics and Epicureans, the high medieval definitions of the human in the Franciscans and Dominicans – the two mendicant Christian orders that staffed the first universities – to the early modern form in which Kant inherited the legacies. (A more elaborate discussion of these parallel streams of thought may be found in Fuller [2011: chap. 2].)

The pedagogical import of these two legacies should not be underestimated. Historically important philosophers can be deemed significant for radically different reasons, which have profound downstream consequences for what is seen as significant in contemporary philosophy. A case in point is René Descartes. To a supply-sider (typically influenced by Anglophone trends), Descartes is someone whose scepticism was born of the potential unreliability of his senses and intellect. In contrast, to a demand-sider (typically influenced by Franco-German trends), Descartes tried to make explicit the special relationship that we have

with God that underwrites the general reliability of our senses and intellect. One consequence of this difference in emphasis is that in the English-speaking world 'epistemology' is naturally aligned with the *philosophy of mind*, which focuses on what happens inside individual heads as we try to secure what little we can know, while in the French and German-speaking worlds epistemology is more naturally aligned with the *philosophy of science*, which focuses of what happens when the structured interaction of individuals produce epistemic wholes, such as a Kuhnian paradigm, that exceed what any of the constitutive individuals could grasp or pursue by themselves.

A good way to think of the overall development of this two-tracked trajectory is in terms of humanity pulled in two directions, up and down – towards the heavens (demand-side, where we re-enact divine creativity) and towards the earth (supply-side, where we re-embed into the natural world). But before Kant's two traditions began to be treated in more explicitly economic terms in the late 19th century, the most natural way to think about their contrasting normative orientations to philosophy had been in terms of the secular professions of law and medicine, specifically *legislation* versus *medication*: On the one hand, the imposition of reason on the world by sovereign will; on the other, the adjustment of the soul to the world by the rationalisation of sentiment. Often this captures the actual preoccupations of the relevant parties – e.g. Plato and Leibniz in law, Aristotle and Locke in medicine.

However, as philosophy came to acquire a distinctly academic cast in the generations after Kant's death, this distinction in life-orientations was domesticated as a division within the ancient discipline of metaphysics into the modern ones of epistemology and ontology. They reflect 19th century developments in Germany (led by Prussia) and Austria, two nations with radically different political outlooks. Whereas Germany aspired to unify modern Europe, Austria struggled to cope with its decline as the seat of Christendom. Translated into philosophy, the German side fully autonomizes epistemology from ontology, arguably rendering ontology a second-order effect of epistemology, i.e. knowledge is *constructive* of being. In terms of medieval theology, the intellect is imposed on the world through the will, as if the human were a deity in the making. In contrast, the Austrian side makes epistemology dependent on, if not a second-order effect of, ontology: i.e. knowledge is *representative* of being. The theological analogue here is that the intellect disciplines the appetites in one's own being, which suggests that humans are the animal that is most adept at self-mastery. If the German world-view moves seamlessly from science to technology

as 'the extensions of man' (Brey, 2000), the Austrian world-view aims to returns thought to the ground of being, which may be defined as 'nature', the 'given', the 'unconscious' – or simply what Freud's and Husserl's philosophy teacher, Franz Brentano called *Evidenz*, which captures the experience of our pre-mediated attachment to reality (Turner, 2010: chap. 6).

While the German and Austrian sides of the divide are both secular, they are secularizing opposing strands of Christianity. The German version secularises from Protestantism and culminates in the collectivization and centralization of knowledge and power, *a la* socialism, while the Austrian view descends from an anti-Enlightenment Catholic backlash that is skeptical of human attempts to approximate divine omniscience and omnipotence; hence, the rise of so-called Austrian economics in the late 19th century. (Joseph de Maistre is an important transitional figure, as noted in Hirschman [1991].) It should come as no surprising that these radically contrasting visions are rooted in strikingly polarised attitudes to David Hume, the philosopher with whom Kant struggled the most during his most creative period. On the one hand, the German idealists saw Humean scepticism as the enemy that had to be overcome through a strongly proactive conception of the mind that distanced the intellect from sensation to impose order on an otherwise indeterminate material world (Beiser, 2000), whereas on the other hand the Austrian realists saw in Hume's scepticism a precautionary check on our intellectual ambitions in terms of what may be adequately justified by experience (Smith, 1994).

Era	Key Philosophical Problem	Rationalism (Legislation)	Empiricism (Medication)
Greek	Form-Matter Relation	Divided (Plato)	Merged (Aristotle)
Roman	Nature of Life	Outworking of Spirit (Stoic)	Coalescence of Matter (Epicurean)
Medieval	Definition of the Human	Apprentice Deity (Franciscan)	Enhanced Animal (Dominican)
Early Modern	Function of Mind	Expression of Reason (Descartes, Leibniz)	Reception of Experience (Locke, Hume)
High Modern	Post-Kantian Division of Metaphysics	Germany (Will as realization of idea in the world): Fichte	Austria (Intellect as adequacy to objects in the world): Brentano

Figure 1. The Two Philosophical Traditions Before Cognitive Economy

In the late 19th century, the 'economic' character of this distinction explicitly came to the fore, with Ernst Mach and Charles Sanders Peirce arriving at some of the most memorable formulations. However, the clearest trace of this transition to 'cognitive economy' transpired in two public talks: W.K. Clifford's 'The Ethics of Belief' (1877) and William James' 'The Will to Believe' (1896), the latter delivered in explicit response to the former. Cast against type, Clifford the mathematician defended a 'supply-side' empiricist epistemology, whereas James the physician backed a 'demand-side' rationalist epistemology. (If 'rationalist' seems like a strange way to cast James, recall the Enlightenment sense of 'Reason = Intellect + Will'.) For the Jamesian voluntary believer, epistemology is about leveraging what we know now into a future we would like to see. In contrast, for the Cliffordian ethical believer, epistemology is about shoring up what we know so that it remains secure as we move into an uncertain future. The former seeks risks and hence errs on the side of overestimating our knowledge, while the latter avoids risk and hence errs on the side of underestimating our knowledge.

In Figure 2, I have marked this version of demand- versus supply-side epistemology in terms of a distinction that emerged in the early modern period of Western philosophy between, respectively, *belief by decision* and *belief by evidence* (Fuller, 2003: chap. 11). In colloquial terms, this is the distinction between providing a 'reason' in terms of the end you are striving to achieve and in terms of the evidence that licenses your claim. In most general philosophical terms, it also captures deduction vis-à-vis induction, as science's *modus operandi*. In the former, one decides upon a hypothesis and submits it to testing; in the latter, one allows the evidence speak for itself without prejudice of prior hypotheses. In the case of 'belief by decision', a decision projects a future from an otherwise indeterminate evidence base through an act of will. Very much like Pascal's 'wager' for the existence of God, to assume an option as one's own is to confirm additional support for its truth. A technological innovation of probability theory was to reduce this process to the assignment of numerical weightings ('degrees of belief') in which the mathematics revealed the commitments one had effectively made. In contrast, 'belief by evidence' envisages evidence as a constraint on an otherwise indeterminate decision procedure by offering the record of experience as the path to follow of least resistance to what lies beyond one's will. An updated version of this mentality from the economics of technology is the idea of 'path-dependency' (Arthur, 1994).

In its day, the distinction between belief by decision and by evidence was seen as a less metaphysically freighted and more psy-

chologically dynamic version of the rationalist's 'innate ideas' and the empiricist's '*tabula rasa'*, respectively. However, probably the most direct historical source for the distinction in the early modern period was the search for a perspicuous way to interpret probability – or, put more poetically, to 'tame chance' (Hacking, 1975, 1990). Should we deal with chance by placing bets with an eye to maximising personal advantage (the standpoint of subjectivist approaches to probability, such as Bayes Theorem) or by registering and adapting to spontaneous emerging tendencies in nature (the standpoint of objectivist approaches to probability, as in normal distribution curves)?

While my own version of social epistemology aims to update James over Clifford, a normally functioning cognitive economy tries to strike a balance between the two positions. For example, in Figure 2, consider countervailing forces of the two 'psychopathologies', adaptive preference formation and confirmation bias. If you are too attracted to novelty, then the weight of the past acts as ballast; whereas if you are instinctively attracted to the familiar, then a mind-set that allows you to see opportunity in novelty is welcomed. The founding sociologist of scientific knowledge David Bloor tapped into this intuition, borrowing (without citation) from Vilfredo Pareto's 'parallelogram of forces' account of ideological formation (Bloor, 1991: chap. 2).

	Belief by Decision (James)	Belief by Evidence (Clifford)
Metaphysics	Transcendentalism	Naturalism
Truth Goal	The Whole Truth (plus some false?)	Only the Truth (minus some half-truth?)
Likely Error	Overestimation	Underestimation
The Nature of Experience	Test of ignorance to be met and overcome	Ground on which knowledge is built
Epistemic Value	Profit (i.e. added value from an investment)	Rent (i.e. derived value from an asset)
The Role of Evidence	Costs (i.e. falsification)	Interest (i.e. confirmation)
Attitude to Risk	Proactionary (risk seeking)	Precautionary (risk averse)
Psychopathology	Adaptive Preference Formation (Too eager to embrace the new)	Confirmation Bias (Too reluctant to reject the old)
Motto	'What doesn't kill me makes me stronger'	'If it ain't broke, don't fix it'

Figure 2. The Two Philosophical Traditions After Cognitive Economy

4. The Problem of the Economic Use of Knowledge Already Produced

Questions remain regarding not only whether resources are used efficiently in the production of knowledge, but also whether the knowledge so produced is used efficiently. About 25 years ago, the University of Chicago library and information scientist Don Swanson (1986), himself originally trained in physics, managed to understand the aetiology of a medical condition simply by reading across literatures in various fields (which the specialists themselves had not done) and piecing together a hypothesis that was then empirically vindicated by a targeted experiment, facilitating the development of an effective treatment. Swanson had been motivated by various bibliometric facts of the sort originally highlighted by the science historian Derek de Solla Price in the 1960s; namely, that an exaggerated version of the Pareto 80/20 statistical principle of elite formation operates in science such that 90% of the citations accrue to 10% of the authors (Price, 1986).

Sociologists have tended to conclude with Robert Merton (1977) that the uncited articles are either truly worthless or their content is eventually incorporated into the cited articles. This has led to institutional incentives for scientists to publish in 'high impact' journals or team up with people whose work is already well cited. Information economists, perhaps drawing on Leibniz's explanation for the presence of evil in (this) the best of all possible worlds, creatively suggest that the mass of relatively uncited work serves to draw attention to the relatively few pieces of work that are well cited – the signal that penetrates the noise, as it were (Dietz and Rogers, 2012). Truth may be known as a whole for all eternity in the divine mind, but time is required for humans to detect it in our necessarily piecemeal fashion; hence the need for the accumulation of experience as registered in the Science Citation Index (SCI). Theologically speaking, a mark of our fallen state is that much effort needs to be exerted in trial and error in order for truth to emerge – but eventually it does for all to see.

Swanson thought of the matter much more straightforwardly. Given the lack of evidence that the uncited articles were actually read, Swanson concluded that they were simply neglected and may well contain valuable knowledge. But this result would require a change in scientific reading habits. Scientists would need to not so strongly focus on the dominant research tendencies in the specific fields where the research was published – in terms of which the uncited pieces no doubt seem irrelevant. Rather, scientists

would have to learn to read across fields to make the connections where the uncited pieces appear as relevant to some other set of problems. An ambitious follow-up to the Swanson result would involve re-deploying research agencies so that they allocate funds to academics who try to solve standing intellectual and social problems by combing and combining the existing literature. These agencies would then commission targeted first-order research aimed at testing knowledge claims the validity of which cannot be agreed simply from a comprehensive and measured reading of that literature. Already "knowledge managers" outside of academia have developed "data mining" procedures for accessing knowledge that, for the most part, academia has failed to exploit but could inspire industrial applications and patents (Fuller, 2002). However, there is no reason why such discovery procedures (or 'retrieval strategies') should remain solely in the private sector and oriented solely toward commercial interests.

One cost-effective policy that library and information professionals could ensure in the name of social epistemology is that, in preparing grant proposals, researchers have identified the full range of precedents for the proposed work, in relation to which the research project would then be formulated. Such a policy would revive the original SCI concern to avoid the duplication of effort in an expanding knowledge system. Given the increasing specialisation of today's researchers, research topics that potentially traverse several disciplinary boundaries may require library and information professionals as co-principals to grant proposals to ensure not only the efficient utilisation of the already available knowledge but also the comprehensive dissemination of the resulting research to relevant academic and non-academic constituencies. This value-added character to the conduct of research is discussed below in terms of 'epistemic justice'.

Were library and information professionals in charge of the knowledge system, no new research into a topic would be commissioned unless the already existing knowledge base had been exploited to its full extent. Thus, resource-intensive methods of original data generation and collection could be replaced, or at least deferred and attenuated, by the development of clever automated search engines ("knowbots") with access to multiple disciplinary literatures. This policy would be very much in the spirit of another University of Chicago librarian, Jesse Shera's (1983) who had coined the phrase 'social epistemology' in the 1960s to keep advances in information technology firmly under the control of the field's original humanist animus. Translated into practice, what Swanson (1986) called "undiscovered public knowledge" sup-

ports the maintenance and use of institutional archives, in the face of increasing budgetary pressures to discard rarely consulted old books, serials and other documents. The general failure of universities and other knowledge-based institutions to follow Swanson's precedent has resulted in an epidemic of "corporate amnesia", aka "mad archive disease" (King, 2002).

But it would be a mistake to conclude that corporate amnesia is merely the by-product of financially motivated negligence. It is also a design feature of science, akin to "planned obsolescence," whereby sciences with more clearly defined and rapidly advancing research frontiers have shorter citation half-lives. In other words, the relevance of each new text to the discipline's current state of play is evaluated quickly, clearly, and irreversibly. This implies a sharpening of the distinction between, so to speak, the discipline's "short term" and "long term" memory, corresponding to a division of labour between a practitioner and a historian of a discipline (Fuller, 2007a: 6–9). Thomas Kuhn (1970) went so far as to argue that the functional differentiation of practitioners and historians of science is itself constitutive of scientific progress, as it operationalises the idea that science moves forward by leaving its past behind. The aforementioned Derek Price, a contemporary of Kuhn's, demonstrated that the harder the science, the sooner most of its literature is consigned to history. "Price's Index" implies that a sense of historicity is automatically generated by new literature falling, as David Hume said of his own first book, "still born from the presses" into oblivion (De Mey, 1982: 120).

Against this backdrop, library and information science stands virtually alone among academic disciplines in its presumption of what might be called a "strong universalism" with regard to knowledge. The field aims to produce knowledge that is "universal" not only in terms of validity but also availability, such that knowledge functions simultaneously as a source of authority and a mode of empowerment. This prospect animates what social epistemologists call "epistemic justice" (Fuller, 2007a: 24–9). Key to the administration of epistemic justice is a reduction in the gap between historian and practitioner knowledge, so as to minimize the power that expertise can exert over lay knowledge. After all, the faster the research frontier recedes from the view, the easier it is for one to be left behind; hence, the familiar phenomenon of a once active researcher who, after a few years in university administration, finds it impossible to return to her original field. This epistemic distance often appears as a layer of new jargon (expressed in both words and symbols) that functions as a barrier to latecomers,

while allowing work to be redescribed as failed, primitive, or incomplete but, in any case, superseded by the new.

Philosophically speaking, a repository for all knowledge would entail access to, as courts demand of witnesses, 'the whole truth and nothing but the truth'. From the standpoint of social epistemology, the field of library and information science exists in the tension between the 'whole' and the 'nothing but' in the slogan, which in Figure 2 we captured in terms of James' and Clifford's views. An expert-driven, discipline-based epistemic culture would have the field focus on nothing but the truth, while a more consumer-driven, democratised epistemic culture would have the field cover truth as a whole. The former strategy is clearly more conservative than the latter, as a focus on nothing but the truth would allow, in statistical jargon, 'false negatives', while a concern for the whole truth would allow 'false positives'. Hanging in the balance is whether library and information science should reproduce the default search patterns of established disciplines. This would run the risk of peremptorily ignoring relevant work, or offer an independent and possibly more adventurous set of recommendations that itself would the risk of throwing up a lot of false leads but may end up, a la Swanson, reorienting more discipline-bound inquirers.

The two main philosophical approaches to social epistemology divide precisely on this point. On the one hand, some see the differentiation of knowledge into distinct expertises as a normal feature of the growth of knowledge. Often this process is depicted in terms of exfoliation or evolution, in both cases implying that expertise is an entitlement earned by those who have trained in and contributed to the discipline historically recognized as authorized to pronounce on a knowledge domain. From this standpoint, library and information professionals identify and police the boundaries separating these knowledge domains, directing users to the expert sources most relevant to their needs. Goldman (1999) revealingly calls this position "epistemic paternalism," implying that an increasingly complex knowledge system requires that users be given increasing guidance on appropriate sources of knowledge. However, it takes for granted that the current division of cognitive labour is itself appropriate and necessary. On the other hand, my own version of social epistemology urges library and information professionals to adopt a more critical stance towards the historically contingent and institutionally entrenched character of existing disciplinary boundaries (Fuller, 1998, 2002).

From this standpoint, Swanson's "undiscovered public knowledge" draws attention to the increasing gaps between domains of

knowledge that result from the tunnel vision induced by disciplinary specialisation. However, this must be distinguished from what the great social science methodologist Donald Campbell (1988) originally called the "fishscale model of omniscience," which implies that personal expertises overlap so that, taken together, there are no epistemic gaps in the community of inquirers. While Campbell's point may describe the aggregate of people's actual knowledge bases, Swanson nevertheless captures people's tendency to interpret what they know of neighbouring fields by the standards of their own fields, thereby limiting the prospects for those fields altering their own frame of reference. Here library and information professionals can facilitate the shifting between disciplinary frames, say, by the design of search engines that cross-classify cognate material so that users are forced to confront items they would not have otherwise deemed relevant to their inquiries. The result would be to shift users into a broader-gauged 'browsing' mode, albeit within the general parameters of their original search. It would strike a small but reliable blow for epistemic justice.

5. Why Did Our Ancestors Think They Knew More Than We Think We Know Now?

To understand the full import of Swanson's achievement, we need to start by recalling that when the Scottish metaphysician James Ferrier introduced 'epistemology' into the English language in the mid-19th century, it was under the influence of German idealism. In particular, he was persuaded by a certain metaphysical interpretation of logic that was originally used to overcome the cognitive impenetrability of Kant's 'noumenon', the realm of things as they are 'in themselves'. The idealists interpreted the 'known' and the 'unknown' as proper subsets of the 'knowable'. In that case, the 'unknowable' makes sense only as a relative concept. In other words, something is unknowable only relative to the specific terms that are used to define what is knowable (Fuller, 2007b: 32–33). William James turned this idealist move into a cornerstone of pragmatism, arguing that certain things are unknowable only because we lack the appropriate 'conceptual scheme' for detecting them. It follows that we should remain open to the prospect of discovering just such a scheme, which would effectively serve as a key that unlocks a previously hidden aspect of reality. James clearly had in mind here psychic

phenomena, the detection of which he took seriously as a scientific project. However, perhaps a more persuasive example was set by James' older contemporary, the chemist Louis Pasteur, a non-conformist Christian who provided a secular update for Augustine's instructions on how to seek God, namely, 'discovery favours the prepared mind'. In Pasteur's case, this amounted to remaining open to the prospect that a solution to a practical problem – namely, spoilage in wine and beer – might require a radical reconceptualisation of the nature of life itself (Stokes, 1997).

Nevertheless, James' insight and Pasteur's example still leave unanswered the exact sense of psychological openness needed for acquiring a new conceptual scheme capable of rendering certain currently unknowable things knowable. But in principle at least, James was suggesting that such things – indeed, any such unknowable things – could be known under the right circumstances. For example, the speed at which the Earth orbits the Sun was knowable only once the Earth was assumed to move, after which the Earth's speed became a matter of routine calculation. Kuhn's theory of scientific revolutions, in which paradigm shifts are likened to the Gestalt switches involved in religious conversions, may be seen as a legitimate heir to this perspective, which Kuhn may have picked up in as a Harvard undergraduate from James' student, C.I. Lewis (Fuller, 2000b: chap. 6). However, the key Kuhnian insight relevant to social epistemology as cognitive economy is that these paradigm shifts may incur transaction costs, as the conceptual scheme of the new paradigm both renders knowable what had been previously unknowable and, more subtly, renders unknowable what had been previously knowable. (This is sometimes called 'Kuhn Loss'.) Thus, when Max Weber spoke of the 'disenchantment' of the world entailed by modern science, he meant *inter alia* that conceptions of purpose in nature that had been so clear to the medieval scholastics became very difficult, if not impossible, to express coherently in the language of pure mechanism (Proctor, 1991: chap. 3). An exemplar of this point is Kant's *Critique of Judgement*, which is best read as just such an act of recovery of a lost sensibility, but in purely modern terms.

An efficient if perhaps surprising way of encapsulating this general idealist-pragmatist construal of epistemology is through the infamous quote about 'known unknowns' and 'unknown unknowns' uttered by US Defence Secretary Donald Rumsfeld during the Iraq War when explaining issues surrounding military strategy. His source for these turns-of-phrase appears to have been Taleb (2007), who had been recently consulted by the Pentagon

(Evans, 2012: chap. 9). The implied logic of this way of understanding cognitive economy defines the realm of the 'knowable' in terms of the matrix presented in Figure 3, which I have adapted to account for the issues of most concern to social epistemology. In what follows, I discuss how a paradigm shift in the knowable in the aftermath of the First World War led people to conclude that, while we have undoubtedly produced more knowledge since the 18th century, we know less of what is knowable than those living, say, a century or more earlier.

	KNOWN	UNKNOWN
KNOWNS	What is published and used	Swanson's 'undiscovered public knowledge'
UNKNOWNS	Experiments where risk is calculable (i.e. social engineering)	Experiments where risk is not calculable (i.e. entrepreneurship)

Figure 3. The Realm of the Knowable

One of the most curious features of modern intellectual history is that educated people today feel that they know much less of all that there is to know than their counterparts did, say, 100 or perhaps even 200 years ago. Clearly the boundaries of the knowable changed dramatically in this period, especially with regard to our framing of the very old, the very large, the very small and the very fast. In addition, humanity's own status as a being uniquely well-positioned to master the knowable has been challenged by people – especially Charles Darwin – in the name of the very 'science' that in the 18th century had been the source of our epistemic empowerment. Nevertheless, from roughly the mid-18th to the early 20th centuries, people thought that they understood – or were on the verge of understanding – the fundamental principles governing natural and human reality, most likely under a unified set of laws. This expectation cut across most scientific, religious, philosophical and political differences. Indeed, one could easily find both 'idealist' and 'materialist' expressions of this faith. This confident organizational approach to inquiry, which in the early 19th century had come to be called 'consilience', was modelled on Newton's grand mathematical physical synthesis of the motions of the heavens and the earth. Indeed, given that mathematics and physics ended up uncovering the Achilles Heel of such confidence, in the form of Einstein and Gödel, it is worth recalling that one of the 18th century's notable mathematicians, Jean d'Alembert, who co-edited the Enlightenment's most influential publishing project, *L'Encyclopédie*, thought of his field as no more

than an adjunct to engineering, dedicated to calculating and measuring entities whose reality had been already vouchsafed by Newton (Collins, 1998: chap. 11). In terms of Figure 3, d'Alembert and his contemporaries clearly thought science dwelled firmly in the realm of 'known unknowns'.

In this context, the main point of empirical research was not to solve ever more specialised academic puzzles but to extend and apply known general principles to contexts where a deep understanding of the case at hand was necessary for the principles to do some palpable good. This point applied no less to social engineering than civil engineering. We would now call it 'policy-based research', and it helps to explain the epistemic orientations of figures as otherwise different in political and moral outlook as G.W.F. Hegel, Auguste Comte, John Stuart Mill, Karl Marx and Herbert Spencer. None of these people founded academic disciplines because they believed that such disciplines were a remnant of medieval scholasticism. This was true even of the one academic in the bunch, Hegel. For him the various disciplines were simply the concrete outworkings of 'consciousness', a secular sense of 'spirit' that Hegel held to be the proper subject matter of philosophy, a 'meta-discipline' that students acquired as the final stage of their self-development, during which they integrated the knowledge they had acquired from the particular disciplines in a personal synthesis that would provide direction for their lives.

All of the above 19th century thinkers are now seen as having underestimated the significance of the new round of disciplinary specialisation that by the end of that century became the hallmark of the modern research university, producing the great mass of 'undiscovered public knowledge', called 'unknown knowns' in Figure 3. Two other 19th century developments stand out here. One is the division of German theology faculties into pastoral and scholarly sides, the latter driven by an indefinite freedom of inquiry, regardless of its implications for matters of faith. To be sure, this *wissenschaftlich* theology had its own radical political consequences, especially in the hands of Ludwig Feuerbach and the 'Young Hegelians'. However, its modus operandi was the popularisation of current esoteric research, not the application of established universal principles (Collins, 1998: chap. 12). A second development, closely associated with William Whewell's coinage of 'scientist' as the name of a specific profession, was the full incorporation of the natural sciences into the university curriculum. These disciplines differed from those of the medieval university in that their knowledge production required mastery of technical skills traditionally associated with the manual arts and where the

primary knowledge output was not a text but an artefact. Despite Whewell's own emphasis on the need for overarching explanations in science, the radical diversification of epistemic practices effectively undermined the drive to integration at the core of the Enlightenment project. This loss of the unifying spirit became self-conscious with the rise of modern library and information science, as discussed in the previous section. It is traceable to the Belgian lawyer Paul Otlet, an inspiration for the logical positivists, who in the early 20th century proposed a universal classification scheme for 'documents' (a broader category than academic writings) to improve the communication of scientific knowledge, even within science itself.

So far all of the above developments in managing the cognitive economy of science were executed in the spirit of the Enlightenment, even in the cases – as we have just seen – where the letter undermined the spirit. The only clear sources of dissent from this general progressive sentiment were the ultra-conservatives (e.g. Joseph de Maistre) who glossed the faith in progress as modernist hubris, presaging a second coming of Adam's Fall. They saw their fears vindicated with the bloody 1789 French Revolution and copycat attempts at violent organized resistance against established authority that punctuated 19th century politics and culminated in the events culminating in the First World War and the Bolshevik Revolution. All of these events appeared to be inspired by humanity's godlike self-belief that it could create anew from first principles societies superior to the ones that they had inherited. These ultra-conservatives accepted the name 'reactionaries' to emphasize that their principled opposition to the progressive tendencies resembled Newton's Third Law of Motion. The reactionaries longed for a return to the Holy Roman Empire, in which an infallible (and inscrutable) Pope, understood as God's emissary on Earth, presides over a heterogeneous domain in which direct control is devolved 'naturally' to the level at which those with the most first-hand knowledge (based on long-standing experience) enjoy the most authority. In today's European Union, this sentiment is codified as the principle of 'subsidiarity' (Siedentop, 2000).

After 1917, former devotees of the Enlightenment began to adopt a secular version of this reactionary perspective in the great march towards today's neo-liberalism. Here the invisible hand of self-organizing markets functioned as the inscrutable deity whose modus operandi was channelled through the principle of subsidiarity (Plehwe and Mirowski, 2009). The *locus classicus* for this metamorphosis is Hayek (1952). Although the position arose as an explicit response to the violence that had been done against

humans in the name of things written in books, its own stance allowed for the violent replacement of books – say, of Marxist or, later, Keynesian macroeconomics – by the personal experience of humans, especially when engaged in free exchange. At this point, it becomes easy to see how the microeconomic interests of shopkeeper capitalism – *la petite bourgeoisie* – might find common cause with the studied irrationalism of *Heimat* ('homeland') thinking promoted by Martin Heidegger (Fuller, 2003: chap. 15 ff). Moreover, we can put a face on this 'missing link' between Austrian free market economics and Heidegger's fundamental ontology, namely, Friedrich von Hayek's Ph.D. supervisor, Othmar Spann, who also served as Max Weber's *bête noire* in his final years (Ringer, 1969: chap. 4). An interesting point of convergence between neo-liberalism's anti-intellectualism towards the economy (i.e. theories are no substitute for experience) and James' voluntarist approach to conceptual schemes is a valorisation of the 'unknown unknowns' quadrant of Figure 3, understood as a sphere of bold social experimentation that I have discussed in terms of 'moral entrepreneurship' (Fuller, 2011: chap. 5; Fuller, 2012: chap. 4). This attitude is core to what in the concluding section I call the 'proactionary' approach to risk (Fuller and Lipinska, 2013).

6. Projecting the Future of Social Epistemology: The Proactionary Imperative

Perhaps the most important overarching problem for social epistemology is the relationship between so-called *moral* and *epistemic* values. Although several different characterizations have been given of this relationship, generally speaking either (1) epistemic values are cast as a special case of moral values or (2) moral values are portrayed as placing constraints on the realization of epistemic values. In the case of (1), epistemic values are envisaged as a kind of 'ethics of belief', again recalling Clifford, which famously defined intellectual discipline as 'belief proportional to evidence'. To be sure, in recent times, a broadened conception of 'epistemic virtue' that harks back to Aristotle and Aquinas rather than Bacon and Mach has taken root in social epistemology, which is more focused on character-based values of the epistemic agent, such as honesty, humility, open-mindedness, tolerance, etc. (Zagzebski, 1996). In the case of (2), epistemic values are portrayed as potentially undermining of the human

condition if they are not pursued within a certain ethical horizon. This orientation conjures up the spectre of the morally indifferent if not inhuman scientist, who in turn requires the oversight of institutional review boards, if not natural law-based restrictions on scientific experiments on human and animals.

As opposed to both of these, my own preferred view involves taking Ockham's Razor to the distinction between moral and epistemic value by arguing that their real difference lies in the time horizon within which a more generic sense of 'value' is expected to be fully realized (Fuller, 2009: chap. 4). Specifically, so-called 'epistemic value' operates with a much longer time horizon for realizing the same sense of 'value' as that of so-called 'moral value'. Here I am identifying 'epistemic value' with the pursuit of truth as an end in itself regardless of the means pursued to achieve it (which in practice amounts to an ethic of efficiency). Given my associating social epistemology both with the original collective teleological project of 'epistemology', and the more recent development of 'post-' and 'trans-' human normative horizons – whereby the values that humans have traditionally tried to achieve come to be realized in some successor 'species' – I have come to believe that we should take seriously the claim of extreme scientists – including Nazi ones – that their research aims to benefit the human condition despite possibly harming many humans in the short-to-medium term.

While we should not give a free pass to scientists who engage in research that places human beings in extreme situations, we also should not pre-empt invalidate their claims by demonizing them as 'pathological', 'inhumane', etc. After all, precedent for the long-termist, 'end justifies the means' ethic of extreme scientists may be found in utilitarian arguments for the welfare of future generations. These arguments would have people discount or deny the value of their own current pleasures in favour of imagined future ones that may well be experienced by others rather than by oneself. Moreover, these arguments may be deployed to justify the systematic redistribution of various resources away from their default users and uses. Thus, one may be felt morally obliged to curb one's personal expenditure of money, carbon, etc. The salient difference between this case and the epistemic value case, I believe, is that the latter is effectively a second-order version of the former. In other words, sacrificing part of the current population to benefit some indefinitely extended future population is like sacrificing a part of one's current self to benefit either a future version of oneself or some future being whose values are sufficiently similar to one's own.

It is only for historical reasons that the relationship between moral and epistemic value has not been seen in this way. In particular, past cases of the dominance of 'epistemic value' (e.g. eugenics) have been coerced rather than freely chosen by those who would be most likely to suffer the immediate consequences. In the emerging world of 'Humanity 2.0' political ideologies, I have characterized the second-order, epistemic value-led option as *proactionary* (suggesting a risk-seeking mentality) and the first-order, moral value-led option as *precautionary* (suggesting a risk-averse mentality). This characterization might be understood as my 21st century way of casting the difference between the demand- and supply- driven epistemologies that has framed the argument of this paper. Both sides require a substantial re-distribution of personal sentiment and material resources. However, the social-epistemic standpoint of the precautionary ideology is that of those living now who then imagine others who would wish to live like them in the future, as opposed to the proactionary ideology, which envisages future life as involving roughly the same degree of dismissal, incorporation and extension of the past as previous generations have done to their predecessors.

REFERENCES

Arthur, W.B. (1994). *Increasing Returns and Path Dependence in the Economy*. Ann Arbor MI: University of Michigan Press.

Beiser, F. (2000). 'The Enlightenment and Idealism'. In K. Ameriks (eds.), *The Cambridge Companion to German Idealism*. (Pp. 18–36) Cambridge UK: Cambridge University Press.

Bloor, D. (1976). *Knowledge and Social Imagery.* London: Routledge.

Brey, P. (2000). 'Technology as extension of human faculties'. In C. Mitcham (ed.), *Metaphysics, Epistemology and Technology, Research in Philosophy and Technology, vol 19*. London: Elsevier/JAI Press.

Campbell, D.T. (1988). *Methodology and Epistemology for Social Science*. Chicago: University of Chicago Press.

Clifford, W.K. (1999). *The Ethics of Belief.* (Orig. 1877). Buffalo N.Y.: Prometheus Books.

Collins, R. (1998). *The Sociology of Philosophies: A Global Theory of Intellectual Change*. Cambridge MA: Harvard University Press.

De Mey, M. (1982). *The Cognitive Paradigm*. Dordrecht: Kluwer.

Dietz, J. and Rogers, J. (2012). 'Meanings and policy implications of "transformative research"'. *Minerva* 50: 21–44.

Elster, J. (1983). *Sour Grapes: Studies in the Subversion of Rationality.* Cambridge UK: Cambridge University Press.

Evans, D. (2012). *Risk Intelligence: How to Live with Uncertainty*. London: Atlantic Books.

Fodor, J. (1981). *Representations.* Cambridge MA: MIT Press.

Fuller, S. (1988). *Social Epistemology*. Bloomington: Indiana University Press.

Fuller, S. (1996). 'Recent Work in Social Epistemology'. *American Philosophical Quarterly* 33: 149-66.

Fuller, S. (2000a). *The Governance of Science*. Milton Keynes UK: Open University Press.

Fuller, S. (2000b). *Thomas Kuhn: A Philosophical History for Our Times*. Chicago: University of Chicago Press.

Fuller, S. (2002). *Knowledge Management Foundations*. Woburn MA: Butterworth-Heinemann.

Fuller, S. (2003). *Kuhn vs Popper: The Struggle for the Soul of Science.* Cambridge UK: Icon.

Fuller, S. (2007a). *The Knowledge Book: Key Concepts in Philosophy, Science, and Culture*. Durham UK and Montreal CA: Acumen Press and McGill-Queens University Press.

Fuller, S. (2007b). *New Frontiers in Science and Technology Studies*. Cambridge UK: Polity Press.

Fuller, S. (2009). *The Sociology of Intellectual Life*. London: Sage.

Fuller, S. (2011). *Humanity 2.0: What It Means to Be Human Past, Present and Future*. London: Palgrave Macmillan.

Fuller, S. (2012). *Preparing for Life in Humanity 2.0*. London: Palgrave Macmillan.

Fuller, S. and Lipinska, V. (2013). *The Proactionary Imperative*. London: Palgrave Macmillan.

Goldman, A. (1999). *Knowledge in a Social World*. Oxford: Oxford University Press.

Hacking, I. (1975). *The Emergence of Probability*. Cambridge UK: Cambridge University Press.

Hacking, I. (1990). *The Taming of Chance*. Cambridge UK: Cambridge University Press.

Hayek, F. (1952). *The Counter-Revolution in Science*. Chicago: University of Chicago Press.

Hirschman, A.O. (1991). *The Rhetoric of Reaction*. Cambridge MA: Harvard University Press.

James, W. (1960). *The Will to Believe, Human Immortality and Other Essays in Popular Philosophy*. (Orig. 1896). New York: Dover.

King, R.G. (2002) 'Mad Archive Disease: Archival Spongiform Encephalopathy, The Loss of Corporate Memory, and the Death of Institutional Archives'. Paper delivered at the combined SSA/CIMA annual meeting, Flagstaff AZ. http://www.homestead.com/infomgmt/files/mad_archive_ disease2.htm

Kuhn, T.S. (1970). *The Structure of Scientific Revolutions*. 2nd ed. (Orig. 1962). Chicago: University of Chicago Press.

Merton, R.K. (1977). *The Sociology of Science*. Chicago: University of Chicago Press.

Plehwe, D. and Mirowski, P., eds. (2009). *The Road from Mount Pelerin: The Making of the Neo-Liberal Thought Collective*. Cambridge MA: Harvard University Press.

Price, D. de S. (1986), *Little Science, Big Science ... and Beyond*. 2nd ed. (Orig. 1963) New York: Columbia University Press.

Proctor, R. (1991). *Value-Free Science?* Cambridge MA: Harvard University Press.

Rescher, N. (1978). *Peirce's Philosophy of Science*. South Bend IN: University of Notre Dame Press.

Ringer, F. (1969). *The Decline of the German Mandarins*. Cambridge MA: Harvard University Press.

Rothschild, E. (2001). *Economic Sentiments: Adam Smith, Condorcet and the Enlightenment.* Cambridge MA: Harvard University Press.

Schaefer, W., ed. (1984). *Finalization in Science*. Dordrecht NL: Reidel.

Shera, J. (1983). "Librarianship and Information Science". In F. Machlup and U. Mansfield, eds., *The Study of Information: Interdisciplinary Messages*. (New York: Wiley), pp. 379–88.

Siedentop, L. (2000). *Democracy in Europe.* London: Penguin.

Smith, B. (1994). Austrian Philosophy: The Legacy of Franz Brentano. La Salle IL: Open Court.

Stokes, D. (1997). Pasteur's Quadrant. Washington DC: Brookings Institution.

Swanson, D. (1986). 'Undiscovered Public Knowledge'. *Library Quarterly* 56 (2): 103-18.

Taleb, N.N. (2007). *The Black Swan: The Impact of Highly Improbable Events*. London: Allen Lane.

Taleb, N.N. (2012). *Antifragile: How to Live in a World We Don't Understand*. London: Allen Lane.

Tetlock, P. (2005). *Expert Political Judgement*. Princeton: Princeton University Press.

Turner, S. (2010). *Explaining the Normative.* Cambridge UK: Polity.

Wissner-Gross, A.D., et al. (2013). 'Causal Entropic Forces'. *Physical Letters* 110: 168702.

Zagzebski, L. (1996). *Virtues of the Mind*. Cambridge UK: Cambridge University Press.