The Mirage of
An Economics of Knowledge*

Philip Mirowski

Version 3.0

May 2007

The author would like to thank the following for their invaluable comments and input: Bruce Caldwell, Wade Hands, Harold Kincaid, Steve Turner, and the audience for the plenary address at the Philosophy of the Social Sciences Roundtable, Tampa FL March 2007. This is a truncated version of an entry for the forthcoming *Handbook on the Philosophical Foundations of Economics.*

Once upon a time, say around the era of David Ricardo and Karl Marx, political economy was primarily concerned with the production of national wealth. This "classical" notion tended to hang on long into the 20th century, well after the invention of neoclassical economics in the 1870s (Mirowski, 1989, chap. 7). Nevertheless, there was no denying that within neoclassical economics, exchange had displaced production as the primary topic of interest; this informed the definition of economics articulated by Lionel Robbins that its proper subject was the "allocation of scarce means among given ends". But subsequently something rather extraordinary happened around the middle of the 20th century, gaining momentum as the century waned. More and more, economics at the cutting edge (as opposed to the textbooks) became relatively cavalier about treating trade as static allocation, and instead became all wrapped up in the image of the market (or the agent) as a processor of information or knowledge. I am not just referring here simply to the phenomenon of the award of the Bank of Sweden Prize to George Akerlof, Michael Spence and Joseph Stiglitz at the dawn of our millennium (Stiglitz, 2002), or to arcane disputes over something called 'common knowledge' in game theory (Aumann, 2000; Geanokoplos, 1992; Samuelson, 2004), or assertions that knowledge was the source of all economic growth (Romer, 1990), or endless arguments over whether rationality is "bounded" or not (Conlisk, 1996). I mean instead that, if your goal in life was to get published in a highly-ranked economics journal, you could no longer safely cast your analysis in terms of the old familiar trope of static allocation. As Kenneth Arrow (in Colander et al, 2004, p.292) put it: "one of the biggest differences between 1950 and 2000 is the much greater role now given to the role of knowledge and information." Economists now bandy about the term 'information' in their papers almost as freely as they had once resorted to the term 'prices'. Clearly something epoch-making had happened to economics, but what precisely was it? Had the *avant-garde* imperiously consigned all that had gone before to a bonfire of the vanities? Sometimes the more impetuous amongst the cognoscenti wrote as though it did:

> There is no single new Law of Economics… The world is not convex; the behavior of the economy cannot be described as if it were solving a (simple) maximization problem; the law of supply and demand has been repealed (Stiglitz, 1985, p.22)

In other instances, various notables hastened to reassure the rest of us that nothing had really changed at all, and that, contrary to most impressions, the info-fascination was just a minor variation on the age-old wisdom of neoclassical economics: (Shapiro & Varian, 1999; Varian, 2002; Kreps, 1997). Both stories appear equally implausible; and this ushers us towards the nub of the present philosophical perplexity. How can it be that everyone seems to believe that there has been some sort of Great Transformation of Economics into a Science of Knowledge, and yet be utterly incapable of producing even a spare consensus on the hallmark doctrines of the New Order?
[i]

Risking ridicule, one can't help but wonder: How do they *know* that knowledge has become central to the discipline? This paradox will lead us into an even deeper question: How is it possible that a neoclassical economic theory, committed to a thoroughly ahistorical non-contextual theory of equilibrium (and notoriously weak on how that equilibrium is attained), could provide an adequate account of the process by which knowledge is gained, interpreted, and understood? At minimum, one might expect that such expert economic epistemologists could inform us how they managed to arrive at such an important breakthrough – but there is no such rendition anywhere to be found. Instead, what we are proffered are a motley of just-so stories about how novel mathematical tools caused the scales to fall from their eyes, or else a fable about how some prophetic scholars woke up to the fact that the neoclassical tradition had been neglecting psychology for over a century, or else a foible of entirely vacuous appeals to the natural progress of science.[ii] Hasn't it become apparent that something is awry with these Wizards of the Knowledge Economy? With each subsequent attempted dismissal of the question, the weasel reveals more of his glinting tooth. Never was there a greater need for philosophical reflection, even in when it comes to a discipline that spurns philosophy as a matter of course.

Here I will put forth the proposition that there can be no such thing as a cogent 'neoclassical' explanation of the recent ambition to propound a neoclassical Economics of Knowledge.[iii] It suffers from the bad habit of presuming what it cannot demonstrate. Upon reflection, the very notion that austere utility maximization could induce economists to develop a naturalized epistemology and a rich cognitive psychology

smacks too much of magic realism. Yet many economists find it perfectly natural to shrug off such paradoxes of self-reference: they have yet to suffer the consequences of their insouciance. To the modern economist, philosophers appear to constitute about as much danger to their prognostications as do literary critics. Therefore, I shall also endeavor to point out that all major existing modern traditions of the Economics of Knowledge have encountered their comeuppance *solely from within*, leading various economists to concede (if not entirely acknowledge) that their own constructions of the epistemology of the agent were structurally incoherent. If I am correct about this, then the widespread contemporary conviction that our science possesses a glittering new toolkit in the form of an Economics of Knowledge is all the more puzzling, and requires explanation in depth.

1. The Three Faces of Knowledge in Postwar Neoclassical Economics

The irreducible diversity and complexity of the postwar orthodoxy turns out to be a necessary prerequisite for understanding the sheer incongruity of the across-the-board watershed in economics in the 20th century from advocates of static allocation to Wizards of Epistemology. Here I would like to take this complexity for granted, by citing some of my earlier work with Wade Hands on the three main schools of postwar American neoclassical microeconomic theory.[iv] Since I am also going to claim there were essentially three conceptual approaches to the economics of knowledge in this period, it might seem as though I will have to visit all 9 cells of the 3x3 possible permutations. Luckily, the actual historical situation was not quite so promiscuous, although neither was it so cut and dried that each school advocated one and only one analytical approach to knowledge. Rather, as happens so often in the history of science (Pickering, 1995; Collins, 2004), the entire virtual space of intellectual possibility was not explored, but neither did everyone converge to a single conceptual option.

It is of paramount importance to observe that no single discrete doctrine served as an all-purpose 'litmus test' for neoclassical orthodoxy in the immediate postwar era. In particular, each of the following 'core' theses was rejected by some substantial subset of card-carrying postwar academic economists in good standing: utility functions exist and

are real, demand curves slope downwards, income effects are significant considerations, it is intrinsic to human nature to maximize something, markets excel in optimizing something, governments are inevitably deleterious to the harmonious functioning of markets, monopoly is detrimental to the successful functioning of markets, supply equals demand, human beings are rational. Rather, interlocking configurations of positions with regard to these and other questions became stabilized into three competing 'schools' in America, whose acolytes were satisfied to vie with one another for adherents and intellectual credibility within an arena of a small set of designated journal outlets, which otherwise served to exclude some other economists as not meeting 'neoclassical' standards. This fostered a situation where neoclassical economics possessed almost no uniform core creed, and yet might begin to expel heterodox elements, starting with self-identified Marxists during the Korean War. It was all the more noteworthy, and served to further obscure matters for outsiders, that two of the three schools were initially physically located at the same geographical location, namely, Chicago. The nascent orthodoxy thus began small and local; but the burgeoning American higher educational system, along with the pervasive abundant military funding of science, served to expand their representation with unprecedented speed.

The numerous ways in which the schools both differentiated their doctrines and sought to co-opt the propositions of their rivals is far too complex to be adequately covered here. Instead, we simply brutally telegraph the tenets most relevant to a history of American microeconomics below:

> A**] Chicago School**. Organized 1946.[v] Initial members were recruited from the wartime "Applied Mathematics Panel", an OR unit located at Columbia and Princeton. Core set: Milton Friedman, Allen Wallis, Aaron Director, George Stigler, and Gary Becker. Demand curves were deemed real and guaranteed to slope downwards, but all the concerns elsewhere about "underlying true determinants" of demand were treated as dispensable. Slutsky equations and integrability conditions were regarded as a minor distraction. Income effects don't matter. Chicago built up (an historically inaccurate) self-account where its origins were purportedly located in Marshallian doctrine. The major determinant of local orthodoxy was the construction of a neo-liberal political doctrine; this dictated strident resistance to Keynesian enthusiasms and denial of the pervasive character of monopoly and market failure. Aggregation problems were treated as irrelevant, and so was imperfect competition. One of their achievements was deployal of a 'workable' empiricism, based upon partial equilibrium analysis and simple ordinary least squares regression analysis. Models

were kept modest and tractable. Major sources of funding are corporations and politically motivated private foundations, particularly the Volker Foundation.

B] **Cowles Commission**. Although Cowles as an organization long predated the neoclassical triumph, its special approach to price theory congealed around new orthodoxy circa 1948 (Mirowski, 2002) when Tjalling Koopmans took over the research directorship from Jacob Marschak. Core set: Koopmans, Marschak, Kenneth Arrow, Gerard Debreu, Leonid Hurwicz, Roy Radner. It was mostly comprised of trained natural scientists moving into economics later in their careers. Cowles was driven from Chicago to Yale in 1954 by rancorous disputes with the Chicago school. In this period it experienced a research shift from earlier attempt to empirically verify Walrasian systems of demand, in favor of treating Walrasian general equilibrium as the Bourbakist 'mother-structure' of all economic theory: this was anointed the Arrow-Debreu model. Slutsky conditions were transformed from empirical proposition needing verification to abstract statement of pervasive interdependence of equilibrium conditions. An alliance with the military (and RAND in particular) strengthened the project of providing an abstract "decision theory" in a context of technology of planning for optimization, and linked economists to the pervasive military organization of science in America. A left-leaning inclination to plan the economy was justified by insisting upon the ubiquity of 'market failures', defined as divergences from Pareto optima. Demand curves were thought to not really exist; only 'demand systems,' which are nearly impossible to empirically verify. The econometric estimation of theory was eventually forsaken in favor of treating the economic agent as a miniature econometrician. Fascination with planning alternative abstract systems of exchange led to the innovation of 'mechanism design' (Lee, 2006). Cowles was also the first of the three schools to seriously engage with game theory (Mirowski, 2002).

C**] MIT**.(Often includes Harvard faculty). Technically founded 1941, effectively 1945. Core set: Paul Samuelson, Robert L. Bishop, Hendrik Houthakker, Robert Dorfman, Robert Solow, George Akerlof, Joseph Stiglitz. A major historiographic problem has been the general lack of awareness of the shape of this school, relative to the others, due to overwhelming presence of Samuelson (1998) in seeking to control his own legacy and portray his *Foundations* as representing orthodox neoclassicism as a whole. The consequence for economists has been to over-rate the importance of the doctrine of revealed preference to the detriment of the larger picture.[vi] 'Revealed preference' purportedly did away with problems of verification of demand curves, which were simply presumed to exist. Problems of inter-related demand, such as complementarity, were played down; integrability conditions reappear as the 'Strong Axiom of Revealed Preference'. This school stood as the major champion of models of 'imperfect competition', which helps explain the relative disinterest in Walrasian general equilibrium. MIT was partial to stories where demand 'did not equal' supply. One major determinant of MIT orthodoxy was the reconciliation of left-liberal statist impulses with imperative to repudiate socialism: Keynesian macroeconomics was seen as the 'middle way'. Yet the distinction between macroeconomics and microeconomics was elided in numerous ways. Partly this

straddle was achieved through an 'ironic' stance towards both the existence of utility and the importance of general equilibrium, which instead promotes 'pragmatic' or toy models (2 agents, 2 goods, 2 states, etc.) illustrating 'principles' which cannot be logically demonstrated as inevitable generic covering laws. Nevertheless, 'Marshallian' partial equilibrium was openly derided.[vii] Empiricism was also treated in a relatively ironic manner: facility with advanced econometric techniques often trumped robust simple estimation. (For instance, Samuelson himself never engaged in any econometric empiricism, but wrote extensively on the agent as econometrician *manqué*.) However, the early success of Samuelson's *Principles* enshrined the MIT approach as the public face of the orthodoxy for decades.

Because the three schools differed so much in their orientation as to what constituted the sanctioned content of legitimate neoclassical price theory, it was almost a foregone conclusion that they would each comprehend the import of the neoliberal notion of the 'marketplace of ideas' in a very divergent fashion. The Chicago department was situated at ground zero of postwar Neoliberalism in America, and hence they were perhaps the quickest to pick up on Friedrich Hayek's notion of the marketplace as ideal information processor and develop it, primarily by scrubbing it clean of any overt Austrian taint, and concertedly seeking to reconcile it with their empirical Marshallian-style portrait of the market. The distance from Hayek's "Economics and Knowledge" [1937] to Friedman's *Capitalism and Freedom* [1962] and Stigler's "Economics of Information" (1961) provides a short but bracing introduction to the rapid domestication of the Neoliberal program within one precinct of American economics. It may not have been exactly what Hayek intended, but by 1960 he was no longer a player in the stabilization of American neoclassical economics. The reaction of Cowles was, by contrast, a cacophony of options concerning how one might begin to upbraid and refute Hayek, and not incidentally, bolster the position of erstwhile Cowles member Oskar Lange, who had portrayed Walrasian general equilibrium as a proof-of-concept of the possibility of central planning. However, the more they felt compelled to engage with Hayek's challenge to their politics, the more they also took on board the conceptual framework of a marketplace of ideas. Of the three postwar schools, MIT was the slowest to take up the gauntlet. It seems possible that much of the delay could be attributed to Paul Samuelson's strident position that one could do scientific neoclassical economics devoid of any commitment to cognitive principles whatsoever; but there is the curious

contrary fact that Samuelson was the first to envision a neoclassical 'science policy', which implicitly presumed certain commitments concerning the treatment of knowledge. Whatever the cause, MIT was soon to join the crowd with its own special version of the cognitive marketplace by the late 1960s.

Now we finally arrive at the 'technical' problem faced by each of the 3 schools in the postwar American context. What would it mean to endow their agent with a capacity for 'knowledge'? After all, one could not seriously dissect the validity of a 'marketplace of ideas' until one registered some commitments as to what knowledge was, and how agents processed it. Here we must insist that *nothing in the previous neoclassical tradition had provided solid dependable guidance as to the appropriate treatment of knowledge in neoclassical economic models*.[viii] There were hints scattered here and there, to be sure, but the urgent imperative to grapple with the 'marketplace of ideas' brought the utter disarray of any epistemic approach in neoclassical economics out into the open. Yet, ultimately, economists did not concoct their ontologisms from scratch, but in each case, depended heavily upon the kindness of strangers, in the person of the external natural sciences,[ix] to suggest the foundations upon which they might erect their mathematizations of knowledge. Out of the myriad of possibilities, postwar neoclassicals settled upon three gross analytical options, which are outlined below in Table 1.

TABLE 1

THREE PARADIGMS OF NEOCLASSICAL 'INFORMATION'

Information is:		
a thing (Shannon)	an inductive index (Blackwell)	symbolic computation (Turing)
Cognition is:		
irrelevant	intuitive statistics & epistemic formal logic	symbol manipulation
Learning is:		
purchase of a commodity	statistical inference	algorithm augmentation
Communication is:		
same as exchange	'signaling'	information transmission

I] Information as a thing/commodity. If a dominant heuristic of postwar economists was "do as little as possible to revise or alter the neoclassical theory handed

down from our forebears" when discussing the operation of the marketplace of ideas, then one can readily appreciate why this option would have initially appeared so attractive. If information was a thing-like object, then it could just be subtended to the commodity space as one more good, and 'nothing' need be changed about the standard maximization model whatsoever. Moreover, a thing-like information absolved the theorists of having to confront whatever model of mind which was supposedly inherent in the utility function. Playing fast and loose with commodity space was a popular pastime in postwar economics – think of the way Gerard Debreu deformed it in order to model uncertainty, or Kelvin Lancaster contorted it to capture 'qualities' – and so editing in 'information' seemed a snap.

The problem immediately arose as to how to 'measure' or 'quantify' this kind of information, and that is where Claude Shannon's "information theory" entered the picture. Shannon had developed an argument which suggested information could be treated just like entropy in physics, comparing it to an enumeration of the number of ways a stochastic microdynamics of symbols could make up a measurable macrostate of messages. Shannon then used the measure to derive theorems about efficient coding procedures to maximize transmission (though telephones, telegraph wires, and other channels) in the presence of noise. I will not reprise the tedious and protracted disputes that broke out in the 1950s and 1960s concerning the attempts by social scientists to co-opt the Shannon definition for their own purposes.[x] A concept fashioned to discuss mechanical obstacles to communication channels may turn out to be utter nonsense when used to discuss the semantics of communication, as many soon came to realize. In most American social sciences, explicit recourse to Shannon information theory disappeared from the theoretical journals by the 1970s; most sophisticated readers of the literature had come to realize that it was not the philosopher's stone, and it wasn't even very impressive as an intuition pump. But that did not exhaust its significance for economics. The Shannon mania of the first two postwar decades had the unintended consequence of bolstering the general impression that scientists could and should treat information as a quantifiable thing, and even as a *commodity*. In practice, it became quite common to conflate the embodiments and encapsulations of knowledge in objects and artifacts as mere epiphenomenal manifestations of a generic 'thing' called information. It was a

reification based largely upon a misapprehension – but it still had untold consequences. One suppressed implication was the bogey of self-reflexivity: "I do not suppose that the information content of this essay could ever be quantified" (Dorfman, 1960, p.585).[xi]

Nevertheless, once knowledge was identified as a 'good', then arguments could begin over just what special sort of good it might be. Perhaps it resembled a capital good, but one capable of metempsychosis, like 'human capital'. Or, perhaps its special conditions of production dictated its status as a 'public good'? Here this version of the economics took a perilous turn, from which it has yet to recover. If you could get people to accept that knowledge was a 'good', it helped if you then began to endow it with all sorts of peculiar qualities. Starting with Samuelson (1954) and Arrow (1962), knowledge was claimed to be a weird sort of thing whose use by one person did not prevent its use by another (in the jargon: 'non-rivalrous'); but also something from which it was intrinsically difficult to prevent another from enjoying the benefits once you bought it (in the jargon: 'non-excludable'). This created all sorts of problem for mathematical modeling, but more to the point, was used in the 1960s-80s to justify state subsidy and provision of this marvelous commodity. But upon the neoliberal turn identified above in Section 2, a curious scholastic argument was subsequently made (Romer, 1990; Warsh, 2006) that the previous characterization had been mistaken, and that knowledge was only 'partially excludable', and distinctively different than 'human capital', rendering it an even more special category beyond 'public goods'. This ontological slipperyness of what, after all, is supposed to be a physical 'given' to the model, is the first symptom of an a more debilitating malady.

II] Information as an inductive index and/or the stochastic object of an epistemic logic

With the development of mathematical statistics, there had been efforts early in the 20th century to link intuitions of a 'good sample' to the amount of 'information' it contained, particularly in the tradition of R.A. Fisher. However, none of these proposals amounted to much outside a small coterie of statisticians. However, in the postwar period, an interesting phenomenon happened where the statistical tools of inductive inference (having spread throughout the social sciences) began to get conflated with models of mind (Gigerenzer & Murray, 1987). Since the story of psychology in the early

20th century consisted of a series of frontal assaults on the conscious mind as executive in charge of rationality, a revanchist movement resorted to the theory of probability to stem the tide.[xii] The situation changed rather radically when mathematical statisticians were brought together with operations researchers and game theorists at the RAND Corporation in the early 1950s. There, especially in the work of David Blackwell, a practice took hold of equating 'information' with measures defined over partitions imposed upon an exhaustive enumeration of states of the world, both actual and virtual.[xiii] Crudely, how much a procedure (it was harder to phrase this in terms of real people) 'knew' about a world was a function of how finely or coarsely it could divide up the possibilities, and thus assign probabilities to eventual outcomes, and the sensitivity with which its detectors could discriminate which of the possibilities had actually obtained. The necessity for game theory to divide and discriminate strategies according to states of the world was an immediate inspiration, but quickly the formalism was rapidly developed in two relatively separate directions: one, as the framework for modern definitions of one version of inductive inference, and the other, as the scaffolding used to assign semantic relations to a modal logic.[xiv] In an alliance with artificial intelligence, it became the basis for formal models of an important class of machine logic.

The formalism of matrices of sharply divided partitions defined over a full enumerations of possible states of the world was an artifact of its origins in the statistics of linear estimators and the matrix algebra of game theory, and so it was essentially a foregone conclusion that postwar economists, many of whom had econometric backgrounds and some of whom had studied game theory, would become intimately acquainted with the state space characterization of 'information' early on. Decision theorists also deemed it a flexible framework within which to couch von Neumann-Morgenstern expected utility. However, there was very little about the formalism that recommended itself as a portrait of *cognition*. Under this description, knowledge was a rather all-or-nothing affair: "Although you may have false beliefs, you cannot know something that is false" (Fagin et al, 1995, p.32). Indeed, it was much better attuned to be recruited as one narrow component of a mechanical inference algorithm, but one where there was no such thing as 'surprises' or unanticipated change, not to mention 'learning', and issues of interpretation were banished. Some took to discussing

mechanical alterations in the underlying partitions as being due to 'signals' emanating from Nature or other agents, but this was a rather cavalier and careless treatment of communication and the implementation of language. Even under such desiccated circumstances, the Blackwell formalism did not prove quite so powerful as some had hoped; for instance, it was shown that a simple one-dimensional index of a certain partition being 'more informative than' another over a given set of states of the possible worlds was generally unavailable, given the partitions imposed only a partial ordering over finite information structures (McGuire and Radner, 1986, pp. 108-119). It was thus a portrayal of knowledge "from the outside", so to speak, preserving more than a tinge of the behavioralist presumptions popular back in the 1950s.[xv] It can be used to discuss a very narrow range of interactions between agents, but only ones where the underlying enumeration of states is completely and comprehensively shared. Here we find the source of the excessive literature in economics journals on the causes and consequences of "common knowledge". For these and other reasons, it would be a mistake to treat the state space formalism as a uniquely credible or comprehensive approach to knowledge. One frequently finds this admitted amongst philosophers and computer scientists, but I have *not once* come across a comparable admission in a text written by an economist:

> We don't feel the semantic model… is the unique 'right' model of knowledge… We do not believe there is a right model of knowledge. (Fagin et al, 1995, p.8)

III] Information as computation

This version of knowledge owes the greatest debt to the postwar development of the computer and the theory of computation, but curiously enough, has proven over time to be the least palatable to the vast majority of neoclassical economists. It predominantly travels under the banner of "computationalism", which tends to identify mental states with the computational states found in (either abstract or tangible) computers (Scheutz, 2002). Computationalism is comprised of many competing visions, ranging from formal symbol manipulation to 'connectionism' to 'machine cognition'; but economists have rarely been sensitive to these controversies within artificial intelligence and cognitive science.

To simplify our exposition, here the processing of information will be equated with symbol manipulation by automata of various computational capacities, with the

Turing Machine occupying the highest rung on the computational hierarchy. The importance of the computational hierarchy is that it permits the proof of impossibility theorems concerning what can and cannot be computed upon machines falling within a particular computational class. Computational approaches have had the prophylactic virtue of ruling out all sorts of physically and mathematically impossible procedures from falling within the purview of a Promethean conception of rationality. Treatment of infinities assumes much heightened significance; implementable algorithms are more highly regarded than in-principle proofs. Furthermore, actual experience with computers have provided all manner of heuristic suggestions as to how to approach cognitive science, perhaps taken to an extreme at certain locations. Indeed, as one Clark Medal recipient has admitted, "if you try and do psychology at MIT, you study computers, not humans" (Matthew Rabin in Colander et al, 2004, p.141).

Early on, the computational metaphor of mind proved a mixed blessing for economists. If one were to seriously entertain the notion of a marketplace of ideas, the problem became where in the economy one would situate the computer. Was each agent a Turing Machine, or perhaps an automaton of less exalted capacity? The von Neumann architecture did seem a bit removed from human cognition, and then there were the interminable disputes of the 1960s-90s over what humans could do that computers could not. Most would admit computers could contain information, but could a computer be seriously thought to be knowledgeable? Or perhaps edging closer to Hayek's vision, the marketplace itself should be treated as one vast Turing Machine, with agents simply plug-compatible peripherals of rather diminished capacities? This problem was compounded by the patrimony of the original neoclassical model, located as it was in non-computable N body mechanics (Mirowski, 1989). Perhaps some aspect of the neoclassical model was shown to be Turing non-computable? The temptation was then to shift the location of the computer to another ontological level in order to evade the unsavory implications. One way to summarize the uneasy love/hate relationship of postwar economists with the computer was that it could not be ignored, but perhaps computer science might be gingerly overlooked. Nevertheless, time and again issues of computability were conceded to be germane to an economics of knowledge (Mirowski, 2002).

****

If economists were poised circa 1950 to become junior epistemologists, whichever of the 3 schools they inhabited, then you might have presumed that they would at least have engaged in a clarifying discussion or two about what they believed knowledge was, and how people achieved it. By and large, this did not happen[xvi] -- hardly a propitious start for our Great Watershed. Instead, many economists boldly set forth confidently asserting all manner of theses about information and the market, but when they were forced to get more specific for the purposes of their mathematical models, they tended to adopt their concepts wholesale from elsewhere, usually locations where serious discussions of information had previously been going on.

2. Vestiges of the Economics of Knowledge

In the interests of philosophical concision (as opposed to fine-grained historical fidelity), we now summarize briefly the trajectories of each postwar school of neoclassical economics through the space of formal treatments of knowledge.

Table 2: The Rough Trajectories of Important Schools of Postwar Economics through Information Space

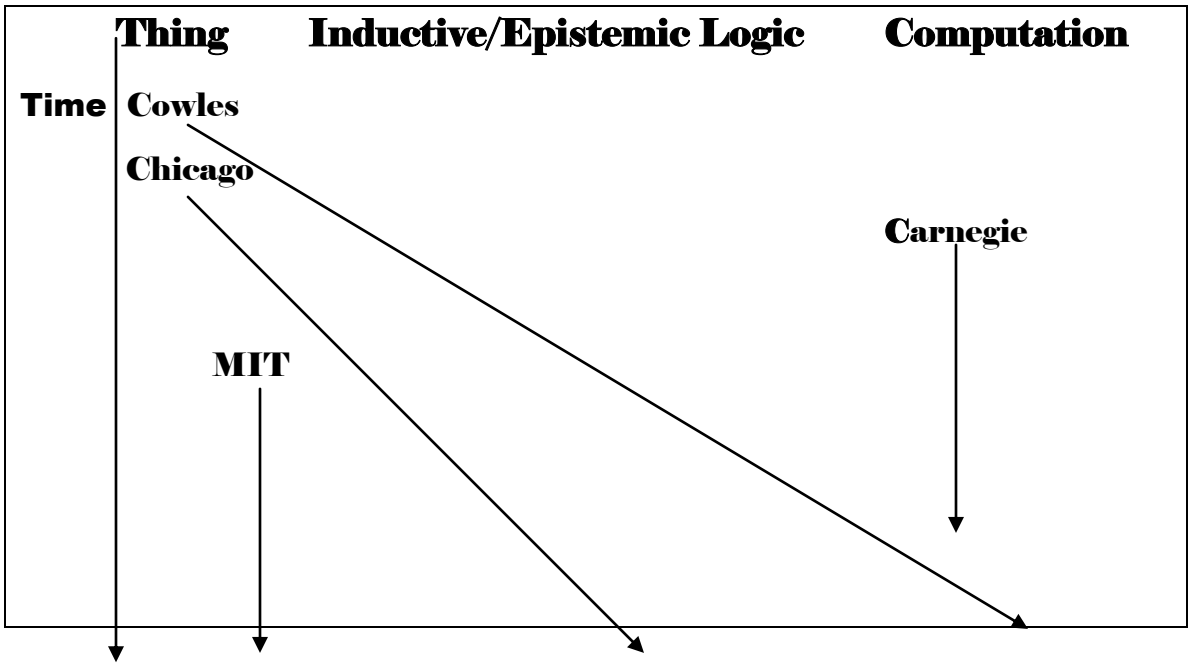

A] The Chicago School. As explained in (Mirowski & van Horn, forthcoming a), Friedrich Hayek was instrumental in consolidating the Chicago School in 1946, and therefore might have been forgiven for expecting that his call for a fundamental reconceptualization of the essence of the market as information conveyance device be pioneered at Chicago; but that is not quite how things turned out. As is well known, Hayek was turned down for a position in the economics department, and had to accept the consolation prize of a professorship at Chicago in the Committee on Social Thought. The early figures of the Chicago school simply presumed that epistemological innovation was superfluous in neoclassical economics, particularly given Friedman's notorious doctrine of 'as if' maximization of utility in his 1953 *Essays in Positive Economics*. Instead, Friedman took the position that simple price theory could already explain knowledge as well as any other commodity, and made no effort to delve into the fine points of epistemology. As some examples of this attitude, his neoliberal instincts told him that the free market of ideas would sanction the privatization of most educational institutions (1962, chap. 6), and that patents simply didn't qualify as illegitimate monopolies (1962, p.127). In this he resembled Gary Becker, who took vague metaphoric appeals to knowledge as 'human capital' and turned them into a protean neoliberal celebration of thing-like knowledge as the pivotal analytical innovation within labor economics. The appeals to knowledge as commodity were even more pervasive in the Chicago Law School, with first Aaron Director (1964) and then Ronald Coase (1974) simply taking the discrete vendibility of ideas as an axiom, and using the putative symmetry between the markets for ideas and commodities as a stick to beat left-leaning academics with over their own lack of consistency. The appeal to "competition" as justification for any set of ideas that neoliberals wished to promulgate was taken to its extreme in the later work of Richard Posner (2005). Even though he was never a faculty member at Chicago, it is our contention that Fritz Machlup (1962, 1980) should be included in this group as well. Chicago also nurtured an attempt to produce neoliberal

sociology of science with roots in the Mertonian tradition, exemplified in the work of Joseph Ben-David (1991, p.11), whom had been informed by Friedman's economics in arguing that state organization of science actually had stunted its development in Europe.

Over time, Chicago began to inch towards a quasi-cognitive approach by linking questions of information to inductive inference. George Stigler entered the lists for the Chicago School in (1961) with an interpretation of cognition as simple statistical sequential search for the lowest price for a given good, which he considered "[his] most important contribution to economic theory" (1988, pp.80-81). His orientation was informed by his experience with the development of sequential analysis at the Applied Mathematics Panel during the war (Wallis, 1980; Klein, 2000), even though he did not avail himself of the full statistical technique. Instead, the model simply posited a stopping rule for visiting a stochastic distribution of 'stores', and then purchasing the item with the lowest realized price known to the agent. This was a characteristically Chicago partial-equilibrium story, since the mere violation of the law of one price would vitiate the microeconomics which Chicago had championed. Full recourse to sequential analysis would not posit a fixed sample size, as Stigler did. Minor improvements along these lines were suggested by Nelson (1970), but by and large, the model did not capture the fancy of the larger profession. One reason might have been the obvious disjuncture between the limited treatment of uncertainty and the larger problem of the role of uncertainty in neoclassical dynamics. It is a Chicago hallmark that this limited form of cognition superimposed upon unchanged static utility maximization neither altered their preferred partial equilibrium analysis in any significant way, nor did it address larger ambitions to treat the market as an information processing device.

The next big Chicago development was a further move in the direction of cognition as inductive inference, but only after a newer generation of neoliberal theorists had suffered a run-in with "serious" cognitive science during their apprenticeship at the fledgling Carnegie School (Sent, 1998, 2002). This was the infamous "rational expectations" movement, which began as a proposition about microeconomics (Muth, 1961), but really caught the attention of economists when it was turned into a weapon to undermine and destroy Keynesian macroeconomics. Figures such as Robert Lucas and Thomas Sargent (1981) adapted the well-established tradition of recasting the cognitive

agent as a little econometrician from decision theory (and Cowles), and augmented it with the stricture that said agent should also use the 'best' theory of price determination in forecasting price movements. It was more than hubris that caused the rational expectations movement to project neoclassical price theory into the recesses of the cranium of the agent: it also permitted the fixed point equilibrium concept to be extended to another area of economics, this time having to do with knowledge. In an unacknowledged bout of reflexivity, neoclassical theory was declared to be the "best" theory of the economy available to the agent (but how did she *know* that?), and therefore also the "best" theory of cognition for that agent. However, there was concomitantly an indirect embrace of Hayek's original neoliberal notion of the market as information processor in the guise of the "efficient markets hypothesis": namely, the proposition that "The Market" managed to incorporate all relevant information into the existing current price, and thus would convey to each and every participant all that he needed to know in order to make a rational economic decision. Yet the divergence from Hayek was also pronounced, since the epistemology of the Chicago-style agent never ventured beyond simplistic models of inductive inference.

We should also take note of a curious development in growth theory (and a retrograde movement back to information as 'thing') which by current criteria should be included under the rubric of the Chicago School (Warsh, 2006). After a hiatus in the 1970s, growth theory was revived in the late 1980s by supposedly dealing with the scandal that 'knowledge' was believed to be the ultimate source of all economic growth, but was treated as unexplained manna from heaven in Solow growth theory. Paul Romer wed the MIT macro fascination with knowledge as a thing with the Chicago advocacy of the commodification of human capital, and produced a model of knowledge as a partially excludable non-rival commodity. For Romer, "knowledge enters production in two distinct ways…[as] a new good that can be used to produce output.. [and which] increases productivity of human capital in the research sector" (1990, p.584). While the mathematical model was a bit awkward, the promise to both 'endogenize' the causes of economic growth and simultaneously to justify the concurrent privatization of intellectual property caused this paper to be one of the all time citation classics in the field of economics.

B] The Cowles School

The advent of information processing at Cowles was a much more complex set of events, and cannot be adequately summarized in a few sentences.[xvii] A crude gloss would point to the fact that Jacob Marschak, Kenneth Arrow and Leonid Hurwicz were all heavily influenced by contemporary developments within cybernetics, but that they all began (under the influence of Claude Shannon and RAND) by treating information as a fungible commodity – "Uncertainty usually creates a still more subtle problem in resource allocation; information becomes a commodity" (Arrow, 1962, p.614) – but only to rather rapidly back off from this option (although never entirely denouncing it), and transfer allegiance to conflating information processing with statistical induction. This move was closely related to their retreat from a full-blown econometric empiricism (for which they had originally gained recognition) in favor of models of the economic agent as himself portrayed as a miniature econometrician. At that point, the original Cowles team fragmented into various unrelated research programs into the economics of information, as the allure of epistemic econometrics palled, and the significance of game theory grew more insistent. Eventually, however, Cowles members in various combinations would explore all three paradigms of information analysis enumerated in the previous section.

There are a number of things to keep in mind about Cowles when observing them foraging about for an economics of knowledge. First, because of their intimate connections with RAND, they were in much closer physical proximity to key natural scientists engaged with innovating new approaches to information than were the other schools. For instance, John von Neumann had made a number of overtures to Cowles economists in the late 1940s (Mirowski, 2002, chap.3). Kenneth Arrow in particular was a close colleague of David Blackwell; Leonid Hurwicz and Stanley Reiter enjoyed close collaborations with various computer scientists. Secondly, many Cowlesmen explicitly admitted that their motivation in the 1940s-60s in discussing information was to refute Hayek (also at Chicago from 1950 onwards), and thus to show that information economics need not have neoliberal implications. Nevertheless, it seems apparent in retrospect that the hunter got captured by the game, in that the frame tale of the

omniscient marketplace of ideas came to dominate much of their own work in mechanism design, asymmetric information, 'failures' of expected utility theory, 'incomplete markets' and a host of other innovations. Finally, it is significant that it was renegade Cowles members (such as Herbert Simon, Alain Lewis, Roy Radner, and Gerald Kramer), and *not* members of the other schools of economics, who came to the realization that tinkering with the utility framework was just too timid a response to the challenge, and struck out to construct a more full-blooded cognitive model, often based upon the Computationalist paradigm, in order to introduce an information paradigm to buttress the edifice of an economics of knowledge.

Some members of Cowles started out believing that the existing Walrasian model was sufficient in and of itself for refutation of Hayek's proposed revision of the marketplace of ideas. Tjalling Koopmans adopted the position that the Walrasian model actually showed that agent cognition was effectively unnecessary, since the individual agent only needed to know his own preferences and parametric prices in order for equilibrium to obtain.

> [O]ne can in particular interpret the proposition as a statement of conditions under which the simplicity of incentive structure and the economies of information handling characteristic of a competitive market organization can be secured without loss of efficiency of allocation… The price system carries to each producer, resource holder, or consumer a summary of information about the production possibilities, resource availabilities and preferences of all other decision makers. Under the conditions postulated, this summary is all that is needed to keep all decision makers reconciled with a Pareto optimal state once it has been established. (Koopmans, 1957, p.53)

The other members of Cowles were not quite so publicly confident that the heritage of Walrasian models adequately addressed this supposed exquisite economy of information: Marschak and Arrow were especially insistent. Much of this discussion within Cowles tried to bundle together the various worries onto the Procrustean Bed of "uncertainty", and one can observe by the mid-1950s Koopmans first floated the trial balloon of blaming this on "missing markets":

> Here, perhaps the most crucial kind of uncertainty…arises from the lack of information on the part of any one decision-maker as to what other decision-makers are doing or deciding to do. It is a puzzling question why there are not more markets

for future delivery through which the relevant information about concurrent decisions could circulate in an anonymous manner.[xviii]

Koopmans was not quite so daunted by these problems as Marschak, about whom "on many occasions during the 1950s and 1960s we heard economists question whether Marschak had not actually left economics for other disciplines, such as psychology [or] information science" (McGuire & Radner, 1986, p.viii). Marschak tried out various paths to his grail of an economics of information (and he was one of the earliest American economists after Hayek to use the term), but none of them seemed to pan out: first he struggled with subjecting Shannon information to a supply/demand framework, and then entertained the Blackwell formalism, only to reject it (Marschak, 1968); he dallied with the idea of transactions costs as capturing informational issues; he also pioneered a computer/organization metaphor which was later to thrive under the rubric of "mechanism design" in the format of his "team theory", which also failed to catch on. He was among the first to participate with professional psychologists in experiments designed to test the limits of decision theory, when that was still an anathema in the economics profession. The failures of Marschak (particularly when compared to the relative success of Arrow) to command the attention of neoclassical economists are perplexing, even in retrospect. Perhaps it had something to do with the fact that the models he explored were embarrassingly progressively removed from the Walrasian general equilibrium that Cowles had come to champion: how to do justice to information in a partial equilibrium framework? Or maybe it had something to do with his level of skepticism over the explanatory power of game theory? Or perhaps it was his relative disaffection from the Bourbakism, which had become all the rage at Cowles?

In any event, it was Kenneth Arrow who became the Cowles poster boy for an economics of information, and indeed, any of the themes covered in this paper be found there at one time or another: knowledge as commodity, information as public good, missing markets, cognition as intuitive statistics, tacit knowledge in the guise of learning-by-doing, decision theory as ersatz psychology, the Blackwell formalism, asymmetric information and moral hazard, bounded rationality, complexity theory, and even cognition as computation. If one does not look too comprehensively at his *oeuvre*, one can find some modicum of support for just about any orthodox approach to the economics of information one might care to promote; and this may account for some of

Arrow's popularity within the profession. The irony of this eclecticism is that at one juncture or another, he has repudiated each and every one of them.[xix] The pattern seemed to be that once a particular research line threatened to invalidate some critical aspect of the neoclassical program or other, Arrow would repudiate the research line, and not the primary notion that neoclassical models were the appropriate vehicle to express the primacy of the marketplace of ideas. This may explain some his most recent crotchety statements, such as:

> The idea that people have difficulty computing the system has a long history; you can see it in Veblen, for example. But nothing followed from this insight. Herb Simon was a great apostle of this view. He's a great figure, and his work did lead to a research program, but in my view, it fizzled out…. As I think more about complexity theory, I become more convinced that there is some sense we will never know how the economy operates. (Arrow in Colander et al, 2004, p.293, 298).

In many ways, the less famous Cowlesmen were more inclined to follow down the consequences of the alternative paradigms to their bitter conclusions. Leo Hurwicz sought to incorporate communication into the Walrasian tradition with his initiation of the program of "mechanism design" (Lee, 2006), if only to better define what neoclassicals meant by the "decentralization of information". Roy Radner sought to ponder even more seriously the implications of cognitive science for the Walrasian program. He explored the consequences of the observation that no agent should be presumed to engage in a trade that depends upon information not available to him at that juncture, and insisted that a Pareto optimum could only be defined relative to a given structure of information, a stricture that, if heeded, would destroy most of the models which today pass under the banner of a neoclassical economics of information. Contradicting Arrow, he insisted that the separation between informational and computational considerations was entirely artificial, and wrote, "The Arrow-Debreu world is strained to the limit by the problem of choice of information. It breaks down completely in the face of limits on the ability of agents to compute optimal strategies" (1968, p.35). Radner's insights have been subsequently ignored for the most part, but will play an important role in the arguments of the next section.

In the modern orthodoxy, the lasting visible heritage of Cowles came with their latching on to the 'state space' formalism as plug-compatible with their general

equilibrium orientation; first pioneered at RAND by David Blackwell (1951), it was now treated in some quarters as the 'standard model' of information in economics (Samuelson, 2004). Although Cowles as an institution left Chicago in 1955, the program it started was continued at RAND, Stanford, Israel, Louvain, and wherever else operations researchers gathered together under military auspices. An important component of OR was the further development of game theory, which was treated as though it were continuous with the neoclassical program.[xx] Game theory was intimately related to the state space formalism and its inductive offshoots, and therefore when strategic cognitive considerations were invoked, it tended to become the paradigm of choice in the treatment of knowledge. However, this rapidly conjured a seeming paradox:

> [I]n the long run, you cannot use information without revealing it; you can use information only to the extent that you are willing to reveal it. A player with private information must choose between not making use of that information – then he doesn't have to reveal it—or making use of it, and then taking the consequences of the other side finding it out… sometimes, in a non-zero-sum situation, you may *want* to pass information to the other side… The question is how to do it so that you can be trusted, or in technical terms, in a way that is incentive-compatible. (Aumann, 2004, p.15)

Notions of strategic mendacity thus entered the pristine marketplace of ideas, like the proverbial serpent in paradise. Robert Aumann was the high priest of this particular theodicy, which attempted to reconcile the ways of Mammon to Man through the instrumentality of what became known as "common knowledge rationality". Aumann's flash of inspiration, which came in the early 1970s at Stanford (Aumann, 2004, p.19), was that the inductive inference version of Blackwell's formalism dictated that if the probabilities of two people for a particular event are common knowledge, then they must be equal. Combine this with the epistemic logic strand of Blackwell, and one arrived at the notion that it would be irrational for two agents to disagree in a fully effective marketplace of ideas. Economists had been slinging around notions of asymmetric or 'private' knowledge for decades, but maybe they just had not taken the lessons of cognition to heart. Or, as (Aumann, 2004, p.19) puts it, "Correlated [Nash] equilibrium is nothing more than just common knowledge of rationality, together with common priors." It would be hard to conceive of a more neoliberal doctrine than that – so much for the original Cowles ambition to repudiate Hayek.

The temptation for all three schools throughout the last half-century has been to elide or slur the distinctions between the three paradigms of information, if only to disguise the ease with which they slid from one modeling tradition to another. Cowles was often most guilty of this, given it covered the most landscape in information space in the postwar period, but perhaps precisely for that reason, it equally produced scholars who warned against the promiscuity:

> Information can be measured, knowledge cannot… There is also an important distinction between 'knowledge theory' and 'information theory.' The former refers to partition models of knowledge, the syntax of knowledge, common and mutual knowledge, axiomatics, and so on… On the other hand, 'information theory' deals with information transmission, noisy channels, entropy and so on. Though related, the two are really quite different. (Aumann, 2005, p.89)

C] MIT School

The path to an economics of information trod by the MIT school looks very different than the road taken by their rivals. For instance, it was only in the late 1960s when some of their foot-soldiers began to take up one of the three paradigms of the treatment of information. But the published record hides a more surreptitious innovation, which dates from the immediate postwar period. Paul Samuelson (2004, p.531) has recently admitted that he helped ghostwrite the postwar bible of American science policy, Vannevar Bush's *Science-the Endless Frontier* (1945). One reads therein that,

> Basic research leads to new knowledge. It provides scientific capital… New products and processes do not appear full-grown. They are founded on new principles and new conceptions, which in turn are painstakingly developed by research in the purest realms of science. (1945, p.11)

In retrospect, it seems clear that this so-called "linear model" of innovation was the precursor to Samuelson's celebrated analytical construct of the "public good" (1954), rather than the other way around. Due to a quirk of fate, Arrow (1962) got credit for the idea that there was something characteristically odd about 'knowledge' that prevented the marketplace of ideas from churning out enough of it, but Samuelson and MIT were the true progenitors. Samuelson and his colleagues were casting about for an economic justification for the state to subsidize the postwar production of knowledge – again we observe the imperative to upbraid Hayek and Chicago on the one hand, but also to rationalize (if not openly endorse) the military control of science on the other (Hounshell,

2000) – and the thought they might hitch it to their advocacy of *macroeconomics*. Knowledge was said to be a prerequisite for economic growth, but there were these curious aspects of non-rivalrousness, uncertainty and zero marginal cost which impeded the marketplace from sufficient investment to guarantee the expansion of the system as a whole. Note well that the MIT approach simply begins with the presumption that knowledge is a commodity without any explicit justification or discussion, trots out the production function, and then lets rip. The capstone to the position was Robert Solow's famous (1957) paper on the sources of technical change, which elevated knowledge to the position of the mysterious 'residual' which explained the preponderance of American growth. This combination of linear model, public good, production functions and national macrostatistics became the only game in town for discussing science policy in the postwar period (Godin, 2005), and bureaucratic knowledge economics in particular, or at least until it was unceremoniously dissolved by the neoliberal tide.

This vision of the economics of knowledge as a kind of accessory to activist Keynesianism can help explain why the MIT school only got with the larger program of choosing one of the three paradigms of knowledge once their Keynesian macroeconomics came under severe attack. We can date their entry into the information sweepstakes with the famous "Lemons" paper by George Akerlof (1970). That paper has an MIT-style toy model which was used to argue a different rationale for government intervention, namely, that asymmetric information (as in the used car market) would cause only faulty cars to be offered on the market, because good cars were constrained to be sold at the same (blue book) price. The 'bottom line' that no cars would be sold at all had (of course) no relationship to the real world used car market; but that was not the point of the exercise. As Akerlof argued in his Nobel lecture:

> [T]he study of asymmetric information was the very first step towards a realization of a dream. That dream was the development of a behavioral macroeconomics… The modeling of asymmetric information was to price theory what the modeling of putty-clay, vintage capital and learning by doing had been to growth theory. (2002, pp.411,413).

Characteristically, these models were used to support political intervention in the marketplace, but this is where things got sticky. MIT (plus Harvard) thought it was bequeathing the profession a previously overlooked wealth of hidden information, secrets

and ambiguous actions, a Kabuki tragedy behind every transaction, but when you looked closely at their models, they really were not all that interested in game theory, which was supposed to be the modality of choice when discussing evasion, mendacity, and strategic behavior. In any event, the Aumann wing of Cowles (discussed above) was busy undermining their simplistic notions of 'private' information and 'hidden' action. The MIT minions believed they had been savaging the Hayekian "marketplace of ideas", but all the while they were amazingly oblivious to epistemological issues; all subtlety appeared lost on these wizards of Mass Ave mentalism. In one case, they conflated the rich complicated history of thought about "competition" to a simple matter of information revelation: "Competition among agents has merit solely as a device to extract information optimally. Competition *per se* is worthless" (Holmstrom, 1982). For the most part, the MIT School was noticeably unconcerned with the problems of the curse of dimensionality as a barrier to cognition (strangely, given their ambivalence about full general equilibrium), and were seemingly uninterested in relevant developments in cognitive science only a few buildings away on the MIT campus. When 'neuroeconomics' caught on with a subset of economists at the turn of the millennium, it bore a "not made at MIT" trademark.[xxi] Perhaps this was a symptom of their sustained agnosticism about belief in the existence of utility.

The most prominent representative of the modern MIT school of the economics of information is Joseph Stiglitz. Although he has been trumpeting the advent of a "new economics of information" for years now (as we noted in the introduction), what is most incongruous about his work is the way it simply ignores almost everything else going on in neoclassical economics, as well as the Blackwell and Turing-inspired paradigms. Whatever else one might think of Cowles and Chicago, it is indisputable that they were intermittently entertaining the pros and cons of different conceptions of knowledge and their consequences for neoclassical theory. But MIT has held itself to different standards:

> [I]t seemed to me the most effective way of attacking the paradigm was to keep within the standard framework as much as possible…While there is a single way in which information is perfect, there are an infinite number of ways that information can be imperfect. One of the keys to success was formulating simple models in which the set of relevant information could be fully specified… the use of highly simplified models to help clarify thinking about quite complicated matters (Stiglitz, 2003, pp.613, 583, 577)

What are the lessons that MIT thinks it has derived from this procedure? As Stiglitz summarized in his Nobel lecture:

> When there is no noise, prices convey all information, and there is no incentive to purchase information. But if everybody is uninformed, it clearly pays some individual to become informed. Thus, there does not exist a competitive equilibrium. (2002, p.395)

It would be something noteworthy if Stiglitz or his co-workers had provided a general impossibility theorem, along the lines of Gödel's incompleteness theorem or Turing's computability theorem, but Stiglitz has repudiated Cowles' stress on general equilibrium (2003, p.580, 620), Chicago's resort to transactions costs (p.573), and doesn't even seriously consider the game theorists' versions of strategic cognition. Indeed, it seems rather heroic to derive any general propositions whatsoever from any of his individual toy models. Stiglitz himself admits this in when he is not engaged in wholesale promotion of his information program.[xxii] Instead, it is possible that "simple" models serve mainly to confuse the issues that beset the quest for an economics of information.

Take, for instance, the famous Grossman-Stiglitz model (1980). The text starts out by positing information as a commodity that needs to be arbitraged (p.393), but claims in a footnote (p.397) that the model of knowledge therein is tantamount to the Blackwell formalism (recall the Aumann quote above), and defines its idiosyncratic notion of 'equilibrium' as equivalence of expected utilities of informed and uninformed agents. Of course, "for simplicity" all the agents are posited identical; how this is supposed to relate to any vernacular notions of divergences in knowledge is something MIT has never been forced to address. When Grossman offered his own interpretation of their joint effort, he took the position that the rational expectations model was identical to the approach in Hayek (1945) [we remind the reader of Section 2 above], that their little toy model had refuted it, that "when the efficient markets hypothesis is true and information is costly, competitive markets break down," and that "We are attempting to redefine the Efficient Markets notion, not destroy it" (1989, p.108).

3. Why there is not as yet a credible economics of knowledge

One lesson we can glean from our survey thus far is that the orthodox "economics of knowledge" literature has been far less concerned with developing a clear position on epistemology than it has been with one of two objectives: (a) to reconcile the bequeathed model of the constrained optimization of utility with something a qualified representative of the natural sciences (but not psychology proper) claims is a scientific treatment of information; and (b) to reprimand Hayek and his followers on their vision of the "marketplace of ideas" and thus promote the favored politics of the particular school, while still treating human cognition as though it resembled a neoclassical market. (It is important to appreciate that *adoption of the* ***Marketplace of Ideas*** is *the* crucial analytical move, and not any nominal pro- or anti-capitalist politics of the individual economist in question.) It is our argument that the bulk of the literature we have cited in Section 2 fails on both counts, and this is the root of our contention that there is not yet an economics of knowledge that deserves the name.[xxiii] The culmination of a half-century of development of each of the three paradigms of information has led to a situation where each is revealed individually incompatible with the core neoclassical price theory.

Why are the "Fundamental Theorems of the Economics of Information" conspicuous by their absence? Why, indeed, does the locution "The Economics of Knowledge" strike many as an oxymoron?[xxiv] We can now begin to proffer a systematic answer to those questions. In this section, we start by suggesting some basic philosophical obstacles, but then because economists tend to despise philosophy, we shall point out how they have undermined their own case in the standard orthodox economics literature. In short, we maintain there are fundamental logical obstacles to equipping the neoclassical agent with a consensus technology to take knowledge on board; but although the neoclassicals have themselves discovered this over the last three decades, they remain loathe to admit it.

We shall proceed by first enumerating four generic contradictions that seem to bedevil all three schools of American neoclassical economics. They are 'generic' in the sense that they do not depend upon the particular paradigm of information (as defined above in Table 1) that the economist in question adopts. In one way of thinking, they are the hidden ontological obstacles to a serious economics of knowledge.[xxv] To make them a bit more vivid and less abstract, we shall dub them: (i) the impossibility of having your

cake and eating it too; (ii) the curse of the schizophrenic agent; (iii) the Wizard of Oz effect; and (iv), the broken bootstrap. Once we have gained some appreciation for the ontological obstacles, we can then point to the ways that they have made themselves manifest in actual models found in the orthodox economics literature. We shall accomplish this by running through the tripartite taxonomy once again: information as thing, information as inductive inference, and information processing as computation. Perhaps through this exercise we can render the obscure objects of desire just a bit less obscure, but not in the vain hope that neoclassical economists will proceed to 'fix' their models as a consequence; they have long shown their disdain for any assistance from philosophers or historians. Rather, our ultimate objective is to provoke some serious alternative economic analysis of modern developments that economists have so far underestimated, such as the corporate takeover of the university, the privatization of science, the 'naturalization' of cognition, and globalization of the neoliberal regime of the organization of knowledge.

(i) The Impossibility of Having/Eating Cake.

The ontological maladies begin, as might be expected, with the history of the commitment to the supposedly non-negotiable precept that the agent comes equipped with fixed and 'well-behaved' set of preferences. This, along with a commitment to ubiquitous maximization, is what gets you your union card as a neoclassical economist in good standing. The story thus imparted as you sat at the feet of your graduate microeconomics instructor in pursuit of that certification is that 'we' once believed in utility as a palpable psychological entity, but Paul Samuelson helped us leave that all behind, and now our minimalist Protestant catechism consists of a few axioms of revealed preference, which boil down to the injunction of consistency in choice, perhaps combined with von Neumann-Morgenstern expected utility as a convenient appendage. It matters little for our present purposes whether anyone actually conforms to this catechism in practice;[xxvi] all we need note here is that you must publicly testify your faith in 'rationality' in order to enter the portals unto the Elysian fields of 'ranked' economics journals.

The first ontological contradiction is that the credo effectively militates against the further equipment of this minimalist wispy agent with whatever cognitive capacities

are deemed requisite to 'solve' the problems inherent in models of information (which are the preferred *topoi* of those very same economics journals that enforce the neoclassical creed). Here is where the trouble really begins. The assertion of agency defined as a fixed, invariant and comprehensive set of preferences devoid of all psychological content, combined with the assertion of agency as an arbitrary complement of cognitive mechanisms allowing the processing of information and the alteration of beliefs/ideas, comes close to tantamount to the simultaneous assertion of A and not-A (which may account for much of the modern polymorphous perversity of the neoclassical agent).[xxvii] The practice has become so pervasive yet so unconscious in Late Neoclassicism that it is difficult to get the average practitioner to even see there is a problem, much less acknowledge it.

Since this really is hard for an orthodox economist to appreciate, let us entertain a brief example. It is well-known that most uncontaminated experimental subjects do not play a Nash equilibrium in the (ultimatum) game 'divide the dollar'. One interpretation suggests this means they are 'irrational'. Yet other defenders of Nash such as Ken Binmore (1999) insist that subjects need to repeat the game a few times to 'learn' to be rational. But that is an unavailing defense, since any specification of Nash equilibrium must include both the knowledge and the cognitive learning capacities of the subjects at the outset, and therefore should be 'always already' incorporated in the Nash calculation. Any attempt to introduce 'knowledge' to save a version of 'rationality' innocent of psychology changes the fundamental constitution of the original model of the agent, and therefore is no longer merely based upon uncontaminated invariant preference. As Bruni and Sugden have recently put it:

> Conventional theory describes the behavior of individuals *who know which actions best satisfy their preferences*. The theory abstracts from the processes by which individuals discover how to satisfy their preferences… So, if preference consistency is interpreted merely as a matter of formal rationality, it is hard to explain the context-independence attributed to discovered preferences. (2007, pp.163,170)

Living the contradiction, a modern economist can sneer at psychology as an 'inferior' social science due to its amorphous orthodoxy, and yet simultaneously become enthusiastic about 'behavioral economics' as somehow rectifying the empirical flaws of neoclassical economics. Embracing the void, the economist can point to papers that

mathematize lexicographic orderings or dispense with the axiom of independence in expected utility theory, but turn around and reject papers for publication that might turn such 'alternatives' into the basis for a general model of the economy. Reveling in the heady air of paradox, the economist can become reconciled to 'bounded rationality' emanating from within the Carnegie School by insisting it is nothing more than constrained maximization applied to the very act of constrained maximization itself (because 'costly' optimization leads to sub-optimization).

Of course, one cannot blithely convict the economist of wanting to stuff incompatible principles together into a single orthodox model—lots of sciences seek a similar grail: consider the repeated attempts to reconcile classical and quantum mechanics. The point we are making here instead is merely that the shotgun marriage of thin invariant preferences with thick cognition will never lead to any foundational model of agency, because any revision or alteration of the one will legitimately be deemed arbitrary from the perspective of the other. The neoclassical agent will always be an unstable compromise between invariant preference field and dynamic information processor.

(ii) Curse of the Schizophrenic agent.

The original neoclassical model from 1870-1940 was first and foremost a model of static allocation of physical goods, based upon the metaphor of a mass-point coming to rest in a field of force (Mirowski, 1989). Equilibrium, while not completely specified (with economics lacking the equivalent of a Hamiltonian dynamics) was more or less straightforward: goods moved around between people through the medium of 'exchange' until the maximum of a utility gradient ensued. What the agents thought about the process, if indeed they could be said to think at all, made no difference whatsoever. There was a single index of success or failure of the market: maximum utility. It was a clean causal story, with a sharp separation of agency from the environment.

But then along came the computer, and with it the attendant ambition for an 'economics of information', and things changed. By endowing the agent with some semblance of cognitive capacities, the very notion of equilibrium began to undergo subtle transformation. In effect, the agent now exhibited *two* distinct motives for exchange: conventional allocate efficiency, and the new notion of cognitive equilibration. It became

conceivable that trade could come to a halt before transactors were content with what they 'knew' or believed about their activities. Divergent beliefs or knowledge might potentially result in trades that would *not* have occurred under the simple regime of static allocation. This was an obvious implication of changing the analytical purpose of the market to be conceived as an information processor.

The disturbing aspect of these developments was that the whole class of phenomena roughly characterized as 'problems with information' was not simply being superimposed as a second-order improvement upon the unchanged core neoclassical orthodoxy (as in the burgeoning literature on 'asymmetric information' 'risk analysis' and so forth), but those problems were tending to subvert the coherence of the foundations of the theory of demand. This prospect really only began to loom large with the displacement of the Walrasian model by game theory as the prime mathematical technique of choice of the orthodoxy.[xxviii] The illusion of continuity within the orthodoxy had been fostered largely as a function of the allegiance pledged by the 1980s to the Nash equilibrium concept, which maintained a foundation in previous utility theory, as well as the idiom of constrained maximization. However, the Nash equilibrium shifted the emphasis to cognitive concerns of what the agent knew about others and their motives, including what the rival knew about what the protagonist knew, and *ad infinitum* … something generally abjured in the prior Walrasian tradition. A truly prodigious literature arose concerning the 'true' meaning of 'common knowledge,' a prerequisite of Nash equilibrium that blurred the boundaries between individual and collective cognition. The position championed by the followers of Herbert Simon at Carnegie Mellon that all the above models only dealt with 'substantive' rather than 'procedural' rationality,[xxix] only served to exacerbate the problem of founding a dynamics upon the basic static neoclassical model.

The simple ontological point to be made is that once the neoclassical agent was endowed with some epistemic abilities, then he/it now came equipped with *at least two separate motives* for exchange, and that these motives need not reinforce one another, but in general, they might instead conflict. This did not bode well for the neoclassical program, which had invested its tough-minded prescriptivism and integrity in their being only a single version of equilibrium (and one that was unique and stable to boot, although

these qualities proved elusive) to which market prices were thought to converge. Our survey of the three paradigms of information below reveals how equilibrium notions fragmented, once caught in a pincers between the dual motives for exchange. Even Hayek never adequately confronted the possibility that allocation and information might be at odds.

(iii) The Wizard of Oz effect

The key departure for postwar neoclassical theory was to essentially buy into the frame tale of the existence of a marketplace of ideas, even though the details might not end up looking very much like Hayek's version of that catallactic universe. Indeed, the earliest way to upbraid Hayek was to insist that the **Marketplace of Ideas** was identical to the neoclassical market model, and that The Market could allocate ideas in the same manner that it allocated widgets. This led to all manner of strange claims being made that neoclassical general equilibrium theory had managed to demonstrate that The Market was the most efficient and parsimonious mechanism in terms of information usage relative to all alternative possible mechanisms of resource allocation.[xxx] For instance, one version of this argument due to Jordan (1982) asserts that competitive equilibrium requires a 'message space' (itself an artifact of the Blackwell formalism) of dimension $n(l-1)$, where n are the number of agents and l is the number of goods, and that any other mechanism requires a message space of higher dimensionality. In a sense, this was a reformulation of Koopmans' original argument at Cowles. Leaving aside quibbles over whether this dimension captures anything of real economic significance, arguments of this ilk are entirely misleading.

The trick to such arguments is to deal entirely with models of static equilibrium, and then paint the neoclassical model as a wonderful embodiment of the parsimonious marketplace of ideas. This conveniently ignores the fact that there is (still) no general theory of dynamic convergence to equilibrium, either for Walrasian general equilibrium or Nash game theory. What should concern economists within this tradition is not only how informationally demanding the presumed mechanism is when the state of equilibrium obtains, but also 'how much' information is required to get us there in the first place. The fruit of decades of effort along these lines has not been reassuring. For instance, the mathematician Steven Smale (1976) proposed a Global Newton Method of

dynamical adjustment for the Walrasian model that did guarantee stability, but only at the price of a truly prodigious informational requirement. In an important paper, Saari and Simon (1978) asked whether they could find 'locally effective price mechanisms' which use less information than Smale's Global Newton method, and answered in the negative. One way to read the Saari/Simon paper is to suggest that any adjustment process leading to an equilibrium from any arbitrary price vector would require an 'infinite' amount of information in a truly general neoclassical world.

Our purpose here is not to dissect the fine points of models of dynamical price adjustment, but rather to make a basic philosophical point—you cannot paint the marketplace of ideas as a marvelously parsimonious and magnificently efficient model of cognition if you can't even demonstrate mathematically that the internal production of neoclassical market equilibrium does not bear information requirements that outstrip any other known algorithmic process. Strip aside the curtain, and you discover to your dismay that the all-powerful wizard is just as weak and flawed as you and me, and that we had been kept in his thrall by some garden-variety *son et lumière* effects.

(iv) The Broken Bootstrap

Suppose, for the sake of argument, that the marketplace of ideas actually manages to price information in such a way that all relevant considerations are somehow embodied in the price data it emits. But then, for the neoclassical economist, the market must be essentially indistinguishable from the cognitive processes people use to process information, although it may differ in details. If that were true, then what precisely is it that induces agents to resort to the market to conduct their cognitive processing rather than just doing it all themselves? Why not 'outsource' most cognition to the mighty **Marketplace of Ideas**? Convenience might be the convenient answer: the market is just *cheaper and easier* than sitting down to think things through from first principles. Sometimes this is phrased using the terminology of 'transactions costs'. But that way lies Bedlam, not Enlightenment.

There are at least two contradictions that arise from this line of reasoning. The first is that, no doubt, we do rely upon others for all manner of assistance, and perhaps even the kindness of strangers, when it comes to accumulating knowledge, but that is because epistemology is inescapably *social or communal*, and not due to the efficacy of

any market phenomena as such (Kusch, 2002). The neoclassical model rules out social epistemology as a matter of course by the way it specifies agency,[xxxi] only to reintroduce it as a *deus ex machina* to extricate itself from paradoxes of positing the neoclassical marketplace of ideas as a superior information processor. This is clearly illegitimate.

The second contradiction arises from paradoxes of self-reference. The posit of 'transactions costs' clearly implies the existence of a 'meta-market' which can set the prices of various formats of market exchange, but that easily leads to an infinite regress. Who sets the prices of the prices of resort to the market? When we shift to an information notion, the paradoxes become more insistent. The market as information processor must itself be priced for us to think about it, but are those prices set 'within' the very same market, or are they banished to some meta-epistemic sphere? In the same way that the Cantor diagonal argument leads to formally undecidable propositions, the price of the marketplace of ideas leads to formally undecidable market prices.

Now, contradictions (i-iv) above partake of the character of in principle objections, the sort that a philosopher might propose. Postwar neoclassical economists have not had much time for philosophical argument (think of Samuelson sneering that real economists do economics, while methodologists just chatter), so one would not expect them to be daunted by anything so flimsy as a mere in principle contradiction. That is why it becomes important to round out this section by demonstrating that each of the three paradigms of the treatment of information in postwar economics have individually come to grief over the postwar period in a purely internalist sense: that is, reputable orthodox economists using orthodox mathematical models have played out the contradictions in their own programs, resulting in the fact that by the millennium no version of the 'economics of information' has emerged unscathed.[xxxii] The landscape, far from being crowded with monumental theorems and general models, is merely dotted with abandoned half-finished shells. This is the effective content of the claim that there is (as yet) no such thing as an economics of information.

Perhaps one of the greatest ironies of the entire situation is that economists have become so very insulated from the culture of which they are nominally a part that they just don't see how absurd the situation looks from outside the walls of their cozy club.

One of the recipients of the 2001 Bank of Sweden Prize given for information economics had the temerity to reveal this insularity in his Prize lecture:

> I was asked recently by a somewhat incredulous questioner (actually a journalist) whether it was true you could be awarded a Nobel Prize in Economics for simply noticing that there are markets in which certain participants do not know certain things that others in the market do know. I thought it was pretty funny. (Spence, 2002, p.435)

Let us decide for ourselves whether the history of the Economics of Information has been a laughing matter. To do so, let us revert back to our Historical Map in Table 2.

4. The Self-refutation of an Economics of Information

[A] Information as a Thing.

Speaking of the Nobel, one of the many incongruities of the 2001 award was that one of the recipients had essentially deconstructed the legitimacy of the entire project of treating knowledge as a commodity. To reprise the quote from Joseph Stiglitz:

> When there is no noise, prices convey all information, and there is no incentive to purchase information. But if everybody is uninformed, it clearly pays some individual to become informed. Thus, there does not exist a competitive equilibrium. (2002, p.395)

While Stiglitz was unable to provide a general theorem based upon a sufficiently generic model, this nevertheless seems a reasonable restatement of what we have called the 'broken bootstrap'. Stiglitz was honored, it seems, for asserting that "informational efficiency" of the **Marketplace of Ideas** is self-contradictory. His papers were published in all the sanctioned orthodox journals. The lessons he prefers to draw from modeling information as a thing is that markets don't need to clear, that the two fundamental welfare theorems don't hold, that there subsists no law of one price, and that in general, supply does not equal demand. I should think the rest of the profession has drawn the rather more obvious conclusion, that one should cease and desist treating information as a simple commodity if one wished to remain a neoclassical economist in good standing.

There was one further reason to reject the thingification of knowledge, at least if one were a partisan of the Cowles School. The more that one declared adherence to Nash game theory, the more it became apparent that the construct of knowledge as a thing you

could hoard and keep private was becoming implausible, since your implacably rational opponent could supposedly divine what you did and did not know from your observable market behavior. From the strategic viewpoint, there were no distinct individuals with their little discrete private bits of fungible knowledge tucked away in the recesses of their craniums populating the marketplace of ideas; there was just common knowledge, and the return of the Germanic Group Mind, where no one possessed the capacity to agree to disagree. This may be one reason why Cowles explicitly repudiated the Shannon formalism by the 1970s. MIT, never really much enamored of game theory, never quite felt the full force of this objection.

Nevertheless, the thinglike conception of knowledge was still quite prevalent in the popular culture, especially in an era of ever-fortified intellectual property rights, so there were a few economists, primarily in business schools and science policy units, who sought to find some accommodation between neoclassical theory and the Thing paradigm. Paul David, its main representative, has insisted that, "Acknowledging the peculiar character of information as an economic commodity is [the] necessary point of departure…" (David, 2003, p.1). The single most common characteristic of this group was a fascination with tacit knowledge.[xxxiii] The idea here was that only some knowledge was 'codifiable', as they liked to put it; the rest was intangible, and passed along outside of marketplace interactions. These economists seemed to believe the tacit/codifiable dichotomy would constitute an escape from the broken bootstrap, or at least the Stiglitz version, since some information just couldn't be purchased. However, (David & Dasgupta, 1994) still argued that maximization held sway, since the problem then became one of choosing the 'optimal mix' of tacit and codifiable information in any given circumstance. This argument became quite popular in the last decade as a neoliberal defense of the privatization of the university and the enclosure of the information commons, even though David later admitted, "we cannot really hope to derive either theoretical propositions or empirical measures regarding whether of not the relative size of the codified portion [of knowledge-P.M.] must be secularly increasing or decreasing" (Cowan et al, 2000, p.230).

I agree with (Nightingale, 2003) that this only appears to rescue the neoclassical tradition by further undermining it. Seriously entertaining the tacit character of

knowledge raises its embodied, social and procedural aspects to the fore, and has been more amenable to constructivist rather than objectivist approaches. It therefore exacerbates the 'cake' contradiction explained above. It is therefore indicative that there are no serious formal models of tacit knowledge within the contemporary neoclassical tradition.

[B] Inductive Inference/ Epistemic Logic

The most telling instance of the philosophical contradiction of the schizophrenic agent has come with the internal development of the Blackwell conception of knowledge, which has been the mathematical economist's epistemic model of choice. It took only a little while after the first specification of the concept of common knowledge for theorists to realize that in a situation where traders are risk-averse, have the same priors and the market clears, then it is also common knowledge that a trader's expected monetary gain given her information must be positive in order for her to be willing to trade. Hence the mere fact that one trader has information which induces her to want to trade at the current price would imply other traders should rationally be unwilling to trade with her: she knows something they apparently do not. Once that happened, it became possible to see that, even in the case where conventional gains in static allocation were possible, informational considerations might serve to stymie any trade (Samuelson, 2004, p.369). This was first mooted by a "no-trade theorem" (Milgrom & Stokey, 1982), which has subsequently been broadened substantially. Furthermore, it was quickly demonstrated that the paradox was not due to any quirk of the original model, but was a direct consequence of the state space model of epistemic logic.[xxxiv] It would seem a tragedy for the neoclassical program that the entire market system would freeze up, just because they had augmented their rational agent with some serious epistemic capabilities. Thus it grew unclear whether the shift to an "economics of information" from the previous static allocation paradigm was really functioning to bolster the orthodoxy, or instead further weakening it.

A similar, but equally damning implication came from Fischer Black,[xxxv] the famous progenitor of Black-Scholes theory of the pricing of financial derivatives. In an important paper (Black, 1986) which wonderfully exemplifies the philosophical contradiction of having your cake and eating it too, Black deconstructed the version of

the Efficient Markets Hypothesis which descended from the Rational Expectations tradition. He was aware of the no trade argument that, if all traders were strategically rational in the neoclassical sense, then prices would embody all relevant information, but that no one would voluntarily take the other side of a proffered trade. But, the volume of trades in financial markets suggested that this portrayal of events could not be correct. Indeed, what most insiders believed is that most financial exchanges were populated by a substantial proportion of "noise traders", that is, people who mistake noise (or their own deluded estimations of their abilities) for real information, and execute trades on that basis. But noise trading renders prices less 'informative' than the ideal posited by rational expectations theory; indeed, the participation of noise traders drives prices away from any conventional notions of 'economic fundamentals'. Nevertheless, noise traders perform a valuable function of creating the conditions for a 'thick' market: they keep the smart insiders in business, so they can provide new information to the marketplace. Therefore, there is no way that prices could convey information in the way proponents of the **Marketplace of Ideas** had claimed.

In some respects, Black's "Noise" resembled Stiglitz's "Impossibility", but with a much better sense for how the real world seemed to work. Perry Mehrling (2005, p.239) asserts that Black did not lose faith in equilibrium, but merely circumvented the schizophrenia of individual rationality by presuming the market 'as a whole' was still rational: in other words, Hayekian Neoliberalism without the Hayekian epistemic commitments. But (evoking the broken bootstrap) for that, who needs neoclassical microtheory?

[C] Information as Computationalism.

The contradictions of the first two paradigms of information were made manifest by honored and revered neoclassical economists; the curiosum of this third case is that the proofs of contradiction were due to more obscure figures.[xxxvi] A number of outsiders to the economics profession, from Michael Rabin to Gerald Kramer to Stephen Kleene had realized that, if one modeled the neoclassical agent as a Turing Machine, then it would be possible to show that many aspects of neoclassical economics could be shown to be non-computable in the sense of formal computational theory. The person who made

this case in detail was the student of Kenneth Arrow and RAND mathematician Alain Lewis. In Lewis (1985) he showed that no finite automata could make the choices which Arrow's choice function formalism had presumed he could accomplish: "It is obvious that any choice function C that is not *at least computationally viable* is ina very strong sense *economically irrational*… the choices prescribed by a *computationally nonviable* choice function can only be implemented by computational procedures that do one of two things: either (a) the computation does not halt and fails to converge, or (b) the computation halts at a non-optimal choice" (1985, pp.45-6). In later papers, Lewis extended the indictment to include the infamous fixed point theorems of Walrasian general equilibrium, the convergence of the core to Walrasian equilibrium, Nash equilibrium, Hurwicz allocation mechanisms, and much else besides. This is the best instantiation of what we have called the 'Wizard of Oz effect' that can be found in the archives of the economics profession.

It seems that few have appreciated just how devastating these results are for the entire program of the Economics of Information. If the marketplace of ideas is thought to operate like a computer, and then one insists upon neoclassical economic theory as the correct and appropriate model of the market, then economists are dealing in delusion, since they regularly endow the market with capacities that *no existing computer can or ever has possessed*. Although it is not a popular opinion in the contemporary profession, it seems hard to escape the implication that neoclassical economics and computers just are incompatible. One may wish (as Hayek did) to portray the entire market institution as resembling a computer, but to do so, one must relinquish any commitment to the neoclassical orthodoxy.

5. Conclusion

I suspect that for many, the arguments contained herein may seem Pyrrhic. Everyone seems to believe that knowledge is the key to economic success, and yet our most-developed schools of economic thought have been mired in the most frightful muddles when it comes to modeling knowledge in an economic setting. The implicit moral is that economists may believe they have left philosophy behind with their high-tech methods, but this is nothing but *hubris* born of isolation and a lack of appreciation for how difficult the problems of knowledge and its comprehension really are.

References

Ancori, Bernard. Et al. 2000. "The Economics of Knowledge: the debate about codification and tacit knowledge," *Industrial and Corporate Change*, (9):255-287.
Akerlof, George. 1970. "The Market for Lemons," *Quarterly Journal of Economics*, (84):488-500.
Akerlof, George. 1984. *An Economic Theorist's Book of Tales.* New York: Cambridge University Press.
Akerlof, George. 2002. "Behavioral Macroeconomics and Macroeconomic Behavior," *American Economic Review*, (92):411-433.
Amin, Ash & Cohendet, Patrice. 2004. *Architectures of Knowledge*. Oxford: Oxford University Press.
Antonelli, Cristiano. 2005. "Models of Knowledge and Systems of Governance," *Journal of Institutional Economics*, (1):51-73.
Apple, Michael, ed.2003. *The State and the Politics of Knowledge.* New York: RoutledgeFalmer.
Apple, Michael. 2006. *Educating the Right Way.* 2nd ed. London: Routledge.
Arrow, Kenneth. 1962. "Economic Welfare and the Allocation of Resources for invention," in Richard Nelson, ed., *The Rate and Direction of Inventive Activity*. Princeton: Princeton University Press.
Arrow, Kenneth. 1984. *The Economics of Information.* Vol. 4 of Collected Papers. Cambridge: Harvard University Press.
Arrow, Kenneth. 1996. "The Economics of Information: an exposition," *Empirica* (23):119-128.
Aumann, Robert. 1976. "Agreeing to Disagree," *Annals of Statistics*, (4):1236-9.
Aumann, Robert. 2000. *Collected Papers, vol 1.* Cambridge: MIT Press.
Aumann, Robert. 2004. "Interview with Sergiu Hart," http://econ.tepper.cmu/md/uploads/05m103R.pdf
Aumann, Robert. 2005. "Musings on Information and Knowledge," *Economics Journal Watch*, (2):88-96.
Baars, Bernard. 1986. *The Cognitive Revolution in Psychology.* New York: Guilford Press.
Babe, Robert. 1994. "Information as an Economic Commodity," in Babe, ed., *Information and Communication in Economics.* Boston: Kluwer.
Bartley, W.W. 1990. *Unfathomed Knowledge, Unmeasurable Wealth.* La Salle, Ill: Open Court.
Ben-David, Joseph. 1991. *Scientific Growth.* Berkeley: University of California Press.
Bender, Thomas & Schorske, Carl. Eds. 1997. *American Academic Culture in Transition.* Princeton: Princeton University Press.
Binmore, Ken. 1999. "Why Experiment in Economics?" *Economic Journal* (109): F16-24.
Binmore, Ken. 2005. "Why the Distinction between knowledge and Belief Might Matter," *Economic Journal Watch*, (2):97-104.

Black, Fischer. 1986. "Noise" *Journal of Finance,* (41):529-43.
Blackwell, David. 1951. "Comparison of Experiments," in Jerzy Neyman, ed., *Proceedings of the Second Berkeley Symposium.*
Blackwell, David. 1953. "Equivalent Comparisons of Experiments," *Annals of Mathematical Statistics,* (24):265-72.
Boldrin, Michele & Levine, David. 2005. "The Economics of Ideas and Intellectual Property," *Proceedings of the National Academy of Science,* Jan. 25, (102):1252-1256.
Boyle, James. 2000. "Cruel, Mean or Lavish? Economic Analysis, Price Discrimination and Digital Intellectual Property," *Vanderbilt Law Review,* (53):2007-2039.
Bruni, Luigi & Sugden, Robert. 2007. "The Road not Taken: how psychology was removed from economics," *Economic Journal,* (117):146-73.
Bush, Vannevar. 1945. *Science-the Endless Frontier.* Washington: Government Printing Office.
Caldwell, Bruce. 2004. *Hayek's Challenge.* Chicago: University of Chicago Press.
Camerer, Colin; Loewenstein, George & Prelec, Drazen. 2005. "Neuroeconomics: How Neuroscience can Inform Economics," *Journal of Economic Literature,* (63):9-64.
Capuro, R. & Hjorland, B. 2002. "The Concept of Information," *Annual Review of Information Science and Technology,* (37):343-411.
Carnap, Rudolf. 1947. *Meaning and Necessity.* Chicago: University of Chicago Press.
Coase, Ronald. 1974. "The Market for Goods and the Market for Ideas," *American Economic Review, Papers and Proceedings.* May:384-391.
Colander, David; Ric Holt & J.B. Rosser, eds. 2004. *The Changing Face of Economics.* Ann Arbor: Univ. of Michigan Press.
Collins, Harry. 2004. *Gravity's Shadow.* Chicago: University of Chicago Press.
Conlisk, John. 1996. "Why Bounded Rationality?" *Journal of Economic Literature,* (34):669-700.
Cooke, Philip. 2002. *Knowledge Economies.* London: Routledge.
Costa, Manuel. 1998. *General Equilibrium Analysis and the Theory of Markets.* Cheltenham: Elgar.
Cowan, Robin; David, Paul; & Foray, Dominique. 2000. "The Explicit Economics of Knowledge: Codification and Tacitness," *Industrial and Corporate change,* (9): 211-253.
Crowther-Heyck, Hunter. 2005. *Herbert A. Simon; the bounds of reason.* Baltimore: Johns Hopkins University Press.
Dasgupta, Partha. 1988. "The Welfare Economics of Knowledge Production," *Oxford Review of Economic Policy,* (4):1-12.
David, Paul. 1997. "From Market Magic to Calypso Science Policy," *Research Policy,* (26):229-255.
David, Paul. 1998. "Common Agency Contracting and the Emergence of Open Science Institutions," *American Economic Review, Papers and Proceedings.* (88:2):15-21.
David, Paul. 2003. "The Economic Logic of Open Science and the Balance between Private Property Rights and the Public Domain in Scientific Information," in J.M. Esanu & P.F. Uhlir, eds., *The Role of Scientific and Technical Data and Information in the Public Domain.* Washington: National Academies Press.
David, Paul. 2004a. "Understanding the Emergence of Open Science Institutions: functionalist economics in an historical context," *Industrial and Corporate Change,* (13:4):571-89.

David, Paul. 2004b. “Can open Science be Protected from the Evolving Regime of IPR Protection?” *Journal of Institutional and Theoretical Economics,* (167):1-26.
David, Paul & Dasgupta, Partha. 1994. “Towards a New Economics of Science,” *Research Policy*, (23): 487-521.
Delanty, Gerard. 2001. *Challenging Knowledge: the university in the knowledge society.* Buckingham: Open University Press.
Demsetz, Harold. 1969. “Information and Efficiency: another viewpoint” *Journal of Law an Economics.* (12):1-22
Director, Aaron. 1964. “The Parity of the Economic Market Place,” *Journal of Law and Economics*, (7):1-10.
Dorfman, Robert. 1960. “Operations Research,” *American Economic Review*, (50):575-623.
Drahos, Peter & Braithwaite, John. 2002. *Information Feudalism: Who owns the knowledge economy?* New York: New Press.
Drori, Gili; Meyer, John; Ramirez, Francisco; & Schofer, Eric. 2003. *Science in the Modern World Polity: Institutionalization and Globalization.* Stanford: Stanford University Press.
Economist Intelligence Unit. 2004. *Scattering the Seed of Invention: the globalization of research and development.*
Egidi, Massimo & Marris, Robin, eds. 1992. *Economics, Bounded Rationality, and the Cognitive Revolution.* Aldershot: Elgar.
Eisenberg, Rebecca. 2001. “Bargaining over the Transfer of Research Tools,” pp.223-249 in R. Dreyfuss, H. First, & D. Zimmerman, eds., *Expanding the Boundaries of Intellectual Property.* Oxford: Oxford University Press.
Eisenberg, Rebecca & Nelson, Richard. 2002. “Public vs. Proprietary Science: a fruitful tension?” *Academic Medicine*, (77):1392-99.
Engelbrecht, H. 2005. “ICT Research, the New Economy, and the Evolving Discipline of Economics: back to the future?” *The Information Society*, (21):317-320.
Fagin, R.; Halpern, J.; Moses, Y. & Vardi, M. 1995. *Reasoning about Knowledge.* Cambridge: MIT Press.
Feldman, Maryann & Link, Albert, eds. 2001. *Innovation Policy in the Knowledge-Based Economy.* Boston: Kluwer.
Floridi, Luciano, ed. 2004. *Blackwell Guide to the Philosophy of Computing and Information.* Oxford: Basil Blackwell.
Foray, Dominique. 2004. *The Economics of Knowledge.* Cambridge: MIT Press.
Friedman, Milton. 1962. *Capitalism and Freedom.* Chicago: University of Chicago Press.
Fuller, Steve. 2002. *Knowledge Management Foundations.* London: Butterworth.
Fuller, Steve. 2005. “Knowledge as Product and Property,” in N. Stehr & V. Meja, eds. *Society and Knowledge.* 2nd ed. New Brunswick: Transaction.
Galison, Peter. 1997. *Image and Logic.* Chicago: University of Chicago Press.
Geanakoplos, John. 1992. “Common Knowledge,” *Journal of Economic Perspectives*, (6):53-82.
Georgescu-Roegen, Nicholas. 1975. “The Measure of Information—a critique,” in J.Rose & C.Bilciu, eds. *Modern Trends in Cybernetics and Systems.* (3):187-217. Berlin: Springer Verlag.

Gigerenzer, Gerd & Murray, David. 1987. *Cognition as Intuitive Statistics.* Hillsdale, NJ: Erlbaum.
Giocoli, Nicola. 2003. *Modeling Rational Agents.* Cheltenham: Elgar.
Glimcher, Paul. 2003. *Decisions, Uncertainty and the Brain: the science of neuroeconomics.* Cambridge: MIT Press.
Godin, Benoit. 2005. *Measurement and Statistics on Science and Technology.* London: Routledge.
Grossman, Sanford. 1989. *The Informational Role of Prices.* Cambridge: MIT Press.
Grossman, Sanford & Stiglitz, Joseph. 1980. "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, (70):393-408.
Gruber, Marc. 2005. Review of Foray (2004). *Journal of Evolutionary Economics.* (15):595-8.
Hacohen, Mordachai. 2000. *Karl Popper: the formative years.* New York: Cambridge University Press.
Hands, D. Wade. 2001. *Reflection without Rules.* New York: Cambridge University Press.
Hands, D. Wade. 2006. "(Continued) Reconsideration of Individual Psychology, Rational Choice and Demand Theory: Some Remarks on Three Recent Studies" *Revue de Philosophie Economique*, (13):3-48..
Hands, Wade & Mirowski, Philip. 1998. "Harold Hotelling and the Neoclassical Dream," pp.322-397 in **Economics and Methodology: Crossing Boundaries**, edited by Roger Backhouse, D.Hausman, U.Maki & A. Salanti, London: Macmillan.
Hartwell, R.M. 1995. *History of the Mont Pelerin Society.* Indianapolis: Liberty press.
Harvey, David. 2005. *A Short History of Neoliberalism.* New York: Oxford University Press.
Hayek, Friedrich. 1948. *Individualism and Economic Order.* Chicago: Regnery.
Hayek, Friedrich. 1982. "The Sensory Order after 25 Years," in W. Weimer & D. Palermo, eds. *Cognition and the Symbolic Process.* Hillsdale: Erlbaum.
Hirschleifer, Jack & Riley, John. 1992. *The Analytics of Uncertainty and Information.* New York: Cambridge University Press.
Holmstrom, Bengt. 1982. "Moral Hazard in Teams," *RAND Journal of Economics*, (13):324-340.
Houghton, John & Sheehan, Peter. 2000. *A Primer on the Knowledge Economy.* Melbourne: Victoria Centre for Strategic Studies.
Hounshell, David. 2000. "The Medium is the Message," in T. Hughes & A. Hughes, eds., *Systems, Experts and Computers.* Cambridge: MIT Press.
Ingber, Stanley. 1984. "The Marketplace of Ideas: a legitimating myth," *Duke Law Journal*,
Innocenti, Alessandro. 1995. "Oskar Morgenstern and the Heterodox Possibilities of Game Theory," *European Journal of the History of Economic Thought*, (17):205-27.
Jaffe, Adam & Lerner, Josh. 2004. *Innovation and its Discontents.* Princeton: Princeton University Press.
Jevons, William Stanley. 1905. *The Principles of Science.* 2nd ed. London: Macmillan.
Jordan, J.S. 1982. "The Competitive Allocation Process is Informationally Efficient Uniquely," *Journal of Economic Theory*, (28):1-18.

Kirman, Alan. 2006. "Demand Theory and General Equilibrium: Down the Wrong Road," in (Mirowski & Hands, 2006).
Klaes, Matthias & Sent, Esther-Mirjam. 2005. "A Conceptual History of the Emergence of Bounded Rationality," *History of Political Economy*, (37):27-60.
Klein, Judy. 2000. "Economics for a Client: Statistical Quality Control and Sequential Analysis," in Roger Backhouse & Jeff Biddle, eds., *Towards a History of Applied Economics*, Durham: Duke University Press.
Koopmans, Tjalling. 1957. *Three Essays on the State of Economic Science.* New York: McGraw-Hill. [reprint Augustus Kelley, 1991]
Kreps, David. 1990. *A Course in Microeconomic Theory.* Princeton: Princeton University Press.
Kreps, David. 1997. "Economics—the Current Position," in (Bender & Schorske, 1997).
Kripke, Saul. 1963. "A Semantical Analysis of Modal Logic I," *Zeitschrift fur Mathematische Logik und Grundlagen der Mathematik*, (24):323
Kusch, Martin. 2002. *Knowledge by Agreement.* Oxford: Oxford University Press.
Landes, William & Posner, Richard. 2003. *The Intellectual Structure of Intellectual Property Law.* Cambridge: Harvard University Press.
Lee, Kyu-Sang. 2004. *Rationality, Mind and Machine in the Laboratory*. Notre Dame PhD thesis.
Lee, Kyu-Sang. 2006. "Mechanism Design Theory Embodying an Algorithm-Centered Vision" pp.283-304 in Mirowski & Hands, (2006).
Lessig, Lawrence. 2001. *The Future of Ideas.* New York: Random House.
Lewis, Alain. 1985. "On Effectively Computable Realizations of Choice Functions," *Mathematical Social Sciences*, (10):43-80.
Lucas, Robert & Sargent, Thomas. 1981. *Rational Expectations and Econometric Practice.* Minneapolis: University of Minnesota Press.
Machlup, Fritz. 1962. *The Production and Distribution of Knowledge in the US.* Princeton: Princeton University Press.
Machlup, Fritz. 1980. *Knowledge: Its Creation, Distribution and Economic Significance.* 3 vols. Princeton: Princeton University Press.
Machlup, F. & Machlup, eds. 1983. *The Study of Information.* New York: Wiley.
Macho-Stadler, Ines & Perez-Castrillo, David. 2001. *An Introduction to the Economics of Information.* 2nd ed. Oxford: Oxford University Press.
Makowski, Louis & Ostroy, Joseph. 2001. "Perfect Competition and the Creativity of the Market," *Journal of Economic Literature*, (39):479-535.
Malone, Cheryl & Elichirigoity, Fernando. 2003. "Information as Commodity and Economic Sector," *Journal of the American Society for Information Science and Technology*, (54):512-520.
Marginson, Simon. 2007. "The public/private divide in higher education: a global revision," *Higher Education*, (53):307-333.
Marschak, Jacob. 1954. "Towards and Economic Theory of Organization and Information," in R.M.Thrall et al., eds. *Decision Process.* New York: John Wiley.
Marschak, Jacob. 1968. "Economics of Inquiring, Communicating, Deciding," *American Economic Review*, (58):1-18.
McCloskey, Donald. 1994. *Knowledge and Persuasion in Economics.* New York: Cambridge University Press.

McGuire, C.B. & Radner, Roy, eds. 1986. *Decision and Organization.* Minneapolis: University of Minnesota Press.
Mehrling, Perry. 2005. *Fischer Black and the Revolutionary Idea of Finance.* New York: Wiley.
Milgrom, Paul & Stokey, Nancy. 1982. "Information, Trade and Common Knowledge," *Journal of Economic Theory*, (26):17-27.
Mirowski, Philip. 1989. *More Heat than Light.* New York: Cambridge University Press.
Mirowski, Philip. 1992. "What were von Neumann and Morgenstern Trying to Accomplish?" in E.R. Weintraub, ed. *Toward a History of Game Theory.* Durham: Duke University Press.
Mirowski, Philip. 2002. *Machine Dreams.* New York: Cambridge University Press.
Mirowski, Philip. 2004. *The Effortless Economy of Science?* Durham: Duke University Press.
Mirowski, Philip. Forthcoming. *SciMart: the new economics of science.* Cambridge: Harvard University Press.
Mirowski, Philip & Hands, D. Wade.. 1998. "A Paradox of Budgets," pp. 260-92 in (Morgan & Rutherford, 1998).
Mirowski, Philip & Hands, D. Wade, eds. 2006. *Agreement on Demand.* Durham: Duke University Press.
Mirowski, Philip & Plehwe, Dieter, eds. Forthcoming. *The Making of the Neoliberal Thought Collective.* Cambridge: Harvard University Press.
Mirowski, Philip & Sent, Esther-Mirjam, eds. 2002. *Science Bought and Sold.* Chicago: University of Chicago Press.
Mirowski, Philip & van Horn, Robert. 2005. "The Contract Research Organization and the Commercialization of Science," *Social Studies of Science*, (35):503-48.
Mirowski, Philip & Van Horn, Robert. Forthcoming a. "The Road to a World Made Safe for Corporations: the Rise of the Chicago School," *Modern Intellectual History*
Mirowski, Philip & van Horn, Robert. Forthcoming b. "Neoliberalism" in Ross Emmett, ed., *Elgar Companion to the Chicago School of Economics*, forthcoming.
Mowery, David & Rosenberg, Nathan. 1998. *Paths of Innovation.* New York: Cambridge University Press.
Mowery, David; Nelson, Richard; Sampat, Bhaven; & Arvids Ziedonis. 2004. *Ivory Tower and Industrial Innovation.* Stanford: Stanford University Press.
Muth, John. 1961. "Rational Expectations and the Theory of Price Movements," *Econometrica*, (29):315-355.
Napoli, Philip. 1999. "The Marketplace of Ideas Metaphor in Communications Regulation," *Journal of Communications*, (49): 151-169.
Nelson, Philip. 1970. "Information and Consumer Behavior," *Journal of Political Economy*, (78): 311-29.
Nelson, Richard. 2004. "The Market Economy and the Scientific Commons," *Research Policy*, (33):455-471.
Nightingale, Paul. 2003. "If Nelson and Winter are only half right about tacit knowledge, which half?" *Industrial and Corporate Change*, (12):149-183.
O'Neill, John. 1996. "Who Won the Socialist Calculation Controversy?" *History of Political Thought*, (17):431-442.

O'Neill, John. 2003. "Unified Science as Political Philosophy," *Studies in the History and Philosophy of Science*, (**)
Pels, Dick. 2005. "Mixing Metaphors: Politics or Economics of Knowledge?" in N. Stehr & V. Meja, eds. *Society and Knowledge.* New Brunswick: Transaction.
Peters, John Durham. 2004. "Marketplace of Ideas: History of a Concept," pp. 65-82 in Andrew Calabrese & Colin Sparks, eds., *Towards a Political Economy of Culture*. Lanham: Rowman & Littlefield.
Phlips, Louis. 1988. *The Economics of Imperfect Information.* Cambridge: Cambridge University Press.
Pickering, Andrew. 1995. *The Mangle of Practice.* Chicago: University of Chicago Press.
Porat, Mark. 1977. *The Information Economy.* Washington: Dept. of Commerce.
Posner, Richard. 2005. "Bad news," *New York Times Book Review*, July 31.
Radner, Roy. 1968. "Competitive Equilibrium under Uncertainty," *Econometrica*, (36):31-58.
Reder, Melvin. 1982. "Chicago Economics: Permanence and Change," *Journal of Economic Literature*, (20):1-38.
Riley, John. 2001. "Silver Signals," *Journal of Economic Literature*, (39):432-78.
Romer, Paul. 1990. "Endogenous Technical Change," *Journal of Political Economy*, (98):S71-S102.
Ross, Don. 2005. *Economic Theory and Cognitive Science—Microexplanation.* Cambridge: MIT Press.
Rothschild, Michael. 1973. "Models of Market Organization with Imperfect Information: a survey," *Journal of Political Economy*, (81):1283-1308.
Saari, Donald & Simon, Carl. 1978. "Effective Price Mechanisms," *Econometrica*, (46): 1097-1125.
Samuelson, Paul. 1954. "The Pure Theory of Public Expenditure," *Review of Economics and Statistics*, (36):387-9.
Samuelson, Paul. 2004. "An Interview with Paul Samuelson," *Macroeconomic Dynamics*, (8):519-42.
Samuelson, Larry. 2004. "Modeling Knowledge in Economic Analysis," *Journal of Economic Literature*, (42):367-403.
Scheutz, Matthias. 2002. *Computationalism: new directions.* Cambridge: MIT Press.
Scotchmer, Susan. 2004. *Innovation and Incentives.* Cambridge: MIT Press
Sent, Esther-Mirjam. 1998. *The Evolving Rationality of Rational Expectations.* New York: Cambridge University Press.
Sent, Esther-Mirjam. 2001. "Sent Simulating Simon Simulating Scientists," *Studies in the History and Philosophy of Science A*, (32):479-500.
Sent, Esther-Mirjam. 2002. "How (Not) to Influence People: the contrary tale of John Muth," *History of Political Economy*, (34):291-319.
Sent, Esther-Mirjam. 2006. "The Tricks of the (No-) Trade (Theorem)" pp.305-321 in (Mirowski & Hands, 2006).
Shapiro, Carl & Varian, Hal. 1999. *Information Rules.* Cambridge: Harvard Business School Press.
Shannon, Claude. 1948. "The Mathematical Theory of Communication," *Bell System Technical Journal*, (27):379-423; 623-656.
Shi, Yenfei. 2001. *The Economics of Scientific Knowledge.* Cheltenham: Elgar.

Slaughter, Sheila, & Rhoades, Gary. 2004. *Academic Capitalism and the New Economy.* Baltimore: Johns Hopkins Press.
Solow, Robert. 1957. "Technical Change and the Aggregate Production Function," *Review of Economics and Statistics,* (**):312-20.
Spence, A.M. 1974. "An Economist's View of Information," in C. Canadia, A. Luke & J. Harris, eds. *Annual Review of Information Science and Technology*. (9):57-78.
Spence, A.M. 2002. "Signaling in Retrospect and the Information Structure of Markets," *American Economic Review,* (92):434-459.
Stigler, George. 1961. "The Economics of Information" *Journal of Political Economy*, (69):213-25; reprinted in *The Organization of Industry*. Homewood: Irwin [1968].
Stiglitz, Joseph. 1985. "Information and Economic Analysis," *Economic Journal—Conference Papers* (95):21-41.
Stiglitz, Joseph. 1993. "Reflections on Economics," in Arnold Heertje, ed., *Makers of Modern Economics, vol. 1* Hemel Hempstead: Harvester Wheatsheaf.
Stiglitz, Joseph. 1999. "Knowledge in the Modern Economy" pp.37-57 in Romesh Vaitilingham, ed. *The Economics of the Knowledge Driven Economy.* London: Department of Trade and Industry.
Stiglitz, Joseph. 2000. "The Contributions of the Theory of Information to 20th century Economics," *Quarterly Journal of Economics,* (140):1441-78.
Stiglitz, Joseph. 2002. "Information and the Change in Paradigm in Economics," *American Economic Review*, (92):460-501.
Stiglitz, Joseph. 2003. "Information and the Change in Paradigm in Economics," pp.569-639 in R. Arnott, B. Greenwald, R. Kanbur & B. Nalebuff, eds. *Economics in an Imperfect World*. Cambridge: MIT Press.
Thomsen, Esteban. 1992. *Prices and Knowledge*. New York: Routledge.
Uebel, Thomas. 2000. "Some Scientism, Some Historicism, Some Critics," pp.151-173 in M. Stone & J. Wolff, eds., *The Proper Ambition of Science.* London: Routledge.
Vaitilingham, Romesh, ed. 1999. *The Economics of the Knowledge Driven Economy.* London: Department of Trade and Industry.
Varian, Hal. 2002. "A New Economy with No New Economics," *New York Times,* Web version, January 17.
Wallis, W. Allen. 1980. "The Statistical Research Group," *Journal of the American Statistical Association,* (75):320-30.
Warsh, David. 2006. *Knowledge and the Wealth of Nations.* New York: Norton.
Washburn, Jennifer. 2005. *University, Inc.* New York: Basic Books.
Wible, James. 1998. *The Economics of Science.* London: Routledge.
Wong, Stanley. 1978. *The Foundations of Paul Samuelson's Revealed Preference Theory*. Boston: Routledge and Kegan Paul.
Woolgar, Steve. 2004. "Marketing Ideas," *Economy and Society,* (33):448-462.
Zappia, Carlo. 1996. "The Notion of Private Information in Modern Perspective," *European Journal of the History of Economic Thought,* (3):107-31.

[i] Fifteen year on, Kreps' (1990, p.578fn) warning is still good advice: "The terms of information economics, such as moral hazard, adverse selection, hidden action, hidden information, signaling, screening and so on are used somewhat differently by different authors, so you must keep your eyes open when you see any of these terms in a book or an article… As a consumer of the literature, you should pay less attention to these labels and more to the 'rules of the game' – who knows what when, who does what when." The only codicil one might add is to replace 'consumer of the literature' with 'epistemically challenged member of the economics community', which better captures the repressed paradox.
[ii] One lame excuse will have to suffice here: "Many scientific discussions focus on knowledge, as researchers have become aware of its importance for value creation on the firm level and wealth creation on the societal level. Yet, there is little common understanding about the special economic properties of knowledge…" (Gruber, 2005, p.595).
[iii] This point was admirably raised in (Hands, 2001, chap.7).
[iv] The primary sources are (Mirowski & Hands, 1998; 2006; Hands & Mirowski, 1998). For those seeking a road map through the present argument, simplifying the trajectory of postwar economics through the thickets of knowledge, subsequent footnotes will identify the key points in the current text as numbered propositions. The three-school division of neoclassical market theory enumerated herein is **Proposition 1**.
[v] The actual conditions surrounding the founding of the Chicago School are highly contested. See (Reder, 1982; Mirowski & Van Horn, forthcoming). The differentiation of the schools discussed in (Hands & Mirowski, 1998) possibly downplays the significant rupture in Chicago price theory in 1946.
[vi] While this complaint might be launched against otherwise perceptive commentaries such as (Wong, 1978; Houthakker, 1983), it can equally be seen as a fault of (Mirowski & Hands, 1998). While revealed preference has become the preferred means of both affirming and denying the centrality of utility functions in the American orthodoxy, it did not provide much in the way of heuristics as to the deployal of demand theory. Many of the obstacles to understanding the MIT/Harvard school come from the inaccessibility of primary archives to researchers, by contrast with the Chicago and Cowles situations. This has particularly blocked research into the patrons of the MIT school.
[vii] One non-Samuelsonian quote will have to suffice: Marshall "…was probably without peer in the delicate art of not letting his inadequate theory get too much in the way of his sensible view of reality" (Bishop, 1964, p.35). On the history of the importation of Cowles-style general equilibrium into the MIT graduate curriculum, see the Duncan Foley interview (p.191 in Colander et al, 2004).
[viii] Previous experience with various audiences has convinced me that this must be deemed **Proposition 2**. Too many *ex cathedra* pronouncements, especially by Nobel winners, seem to have driven this fact from the minds of most contemporary trained economists.
[ix] The work of Shannon is surveyed in (Mirowski, 2002, pp. 68-76); the role of Blackwell at RAND is briefly covered in (Mirowski, 2002, pp.379-389) and the contributions of Turing are described in (Mirowski, 2002, pp. 80-88). Actually, the role of John von Neumann in introducing all three paradigms into economics was substantial, and his legacy is the subject of the book cited. The identification of these three particular theories as primary sources for the subsequent evolution of the economics of knowledge within neoclassicism is **Proposition 3**.

[x] On the inappropriateness of the measure for economic and other purposes, see (Tribus in Machlup & Mansfield, 1983; Floridi, 2004, pp. 46-57; Mirowski, 2002, pp.73-76; Arrow in McGuire & Radner, 1986).
[xi] The notion that those devious relativists who thrive in Science Studies are the only cadre who are susceptible to the perils of reflexivity is one of the sillier arguments made by modern philosophers.
[xii] "History has witnessed the attempt to make probability theory coherent with what was believed to be rational thought, and it has seen efforts to reduce rational thought to probability theory. For instance, what was believed to be rational judicial and economic thought actually determined the way in which probability theory developed mathematically" (Gigerenzer & Murray, 1987, p.137).
[xiii] The historical background to this development is covered in (Mirowski, 2002, pp.380-386). A nice introductory analytical treatment from the standpoint of epistemic logic is (Fagin, et al, 1995).

[xiv] In this latter case, we observe one of the few instances where professional philosophers played a significant role in the development of a notion of knowledge that later became important in economics. The reason this happened was that many of the philosophers in question were also active at RAND in their other capacity as operations researchers. The story begins with Rudolf Carnap (1947), and reaches a high level of development with Saul Kripke (1963).

[xv] "Our notion of knowledge [herein] in a multi-agent system is best understood as an external one, ascribed by, say, the system designer to the agents. We do not assume the agents compute their knowledge in any way, nor do we assume they can necessarily answer questions based on their knowledge" (Fagin et al, 1995, p.9). This quote reveals that this state-space conception is far removed from the computational conception described below.

[xvi] (Machlup, 1962) was clearly one attempt, but actually ended up avoiding most of the thorny epistemological issues, as well as missing out on contemporary developments. In any event, it was roundly ignored by theorists.

[xvii] For a detailed summary, see (Mirowski, 2002, pp.370-389). Even there, many important Cowles initiatives are left unexplored.

[xviii] "Comments in Thursday afternoon session" Conference on Expectations, Uncertainty and Business Behavior, Pittsburgh, Oct. 27-29, 1955, Box 5 folder 81, Tjalling Koopmans Papers, Sterling Library, Yale University. Note that, even though Koopmans was close to von Neumann in this era, he did not entertain the notion that game theory was a better formalism for addressing these questions.

[xix] For the repudiation of the Shannon concept, see Arrow in (McGuire & Radner, 1986). For the admission that his models had little to do with cognitive information processing, see (Arrow, 1984, p.200). "There is no general way of defining units of information" (Arrow, 1996, p.120). For Arrow's role in suppressing the work of Alain Lewis, which plays an important role in the next section, see (Mirowski, 2002, pp.427-36).

[xx] In fact, only a small part of game theory was developed in conformity with neoclassical models, but since that subset (primarily involving Nash equilibrium theory) later grew to such dominance in the economics profession, we shall restrict ourselves to that tradition for the purposes of this paper.

[xxi] One might observe this by noting that the survey article on neuroeconomics by (Camerer et al, 2005) might be seen as a laundry list of all the ways in which 'revealed preference theory' is flat out wrong.

[xxii] "Unfortunately, we have not been able to obtain a general proof of any of these propositions. What we have been able to do is analyze an interesting example" (Grossman & Stiglitz, 1980, p.395).

[xxiii] This explicit rendering of the problematic of postwar neoclassical economics is **Proposition 4**.

[xxiv] Here, of course, we refer to the opinions of *outsiders* and not card-carrying members of the economics profession in good standing. Some examples of those we have in mind are (Boyle, 2000; Apple, 2006; Marginson, 2007).

[xxv] The assertion and enumeration of these generic ontological obstacles to a viable economics of knowledge is **Proposition 5**.

[xxvi] Again we point to (Hands, 2006) as a perceptive theological deconstruction of actual events.

[xxvii] This has long been the complaint of philosophers critical of behaviorist psychology, coming from an appreciation for the literature of Continental philosophy. Here we might mention Charles Taylor, Alisdair MacIntyre, Michel Foucault, and Ian Hacking.

[xxviii] See, for instance, Peyton Young: "game theory challenged a basic tenet of classical economics because it called attention to situations in which individuals acting in their own self-interest do not necessarily arrive at a social optimum. Previously these situations had been perceived as exceptional or peripheral; game theory showed that they are ubiquitous" (in Colander et al, 2004).

[xxix] On Simon, see (Sent, 2001; Klaes & Sent, 2005; Crowther-Heyck, 2005). It has been since noted that one culprit in this regard was the Blackwell 'state space' formalism, which effectively banished procedural questions (Samuelson, 2004, p.400).

[xxx] These papers began with the Cowles doctrine of Koopmans cited above in Section 3, and continued with the tradition of mechanism design associated with the names of Hurwicz and Reiter. For further considerations of this literature, see (Costa, 1998; Lee, 2006; Kirman, 2006).

[xxxi] I am aware there are many philosophers who call themselves social epistemologists and yet make direct use of neoclassical models of the agent, such as Philip Kitcher and Alvin Goldman. Since I have critiqued their work in detail elsewhere (2004), I will here merely reiterate that the implications the current critique of economics have more far-reaching consequences than previously allowed.

[xxxii] The tracing of each paradigm of the postwar treatment of information in economics found in Table 1 to its culmination in its own antithesis is **Proposition 6**.

[xxxiii] The concept originated with Michael Polanyi; but these economists rarely read Polanyi, much less acknowledged the subtleties of his position (Mirowski, 2004, ch. 2). For the purposes of this paper, we shall identify Paul David (David & Dasgupta, 1994; Cowan et al, 2000) and Dominique Foray (2004) as representatives of this tendency. Paul Nightingale (2003) provides a wide-ranging critique of this position; see also (Ancori et al, 2000).

[xxxiv] This is explained in (Fagin et al, 1995, p.184). See also (Samuelson, 2004; Sent, 2006).

[xxxv] For historical background on Fischer Black, as well as a pellucid explanation of this paper, see (Merhling, 2005).

[xxxvi] A more elaborate history of these contributions may be found in (Mirowski, 2002, pp.422-436).