STUDIES IN EMERGENT ORDER VOL 2 (2009): 76-108

The Adaptive Systems Theory of Social Orders

Thomas J. McQuade and William N. Butos

1. Introduction

Many social orders can be understood as instances of adaptive systems – complexes of active components whose interactions implement a persistent but mutable network-like structure which is adaptable to the characteristics of its environment to which it is sensitive. This general model of social structure, being neither exclusively reductionist in that it encompasses emergent phenomena and adaptive reactions at the system level, nor exclusively holistic in that it pays due attention to the lower-level interactions which drive the system, can serve as a fertile source of novel ideas for investigating a range of social complexes and attempting to understand the particular architecture of each and its adaptability and stability characteristics when subjected to environmental pressures of various sorts.

The inspiration for this approach is Hayek's (1952) explanation (rediscovered in an expanded and more sophisticated form in today's neural network theory) of how our brains are able to create the array of sensory qualities by which we perceive events. Hayek's central idea was that the mutable network of interacting neurons in the brain functions as a "map" of the previously experienced environment in the sense that it instantiates a classification of the stimuli that have impinged on the neural system from that environment. The interactions generated in the network by current stimuli are a "model" of the system's current environment.[1] The map is built from experience, being modified by strengthening of neuronal connections when new experience confirms old and by the formation and detachment of connections when new experience produces activation patterns different from those previously experienced. The system is, in this very particular way, self-organizing. Its relational classes, instantiated as distinguished patterns of neural activity which are attractors in activity space, comprise the current state of its knowledge of its environment.

Particular social arrangements, if conceived as networks of people interacting via institutionalized transactions which provide both local incentives to interact and global feedback reflecting the systemic consequences of such interaction, are brain-like in a limited but important respect – specifically, the interactions within the system implement a classifying process on stimuli impinging on the system, and this process can induce real changes in interaction patterns that, in turn, engender adaptive reactions of the system as a whole to changes in its environment. Such arrangements are designated here as "adaptive classifying systems".[2]

The claim that many social orders can be usefully understood as adaptive classifying systems requires a three-pronged defense:

1. An enumeration of the general attributes of adaptive systems, showing how these characteristics can be construed as matching relevant aspect of social arrangements.
2. A clear explanation, for each social arrangement of interest, of what is being classified, how is the classification being carried out, what sorts of classes are being generated, how these classes change with experience, what sorts of responses are generated, and how the effects of the responses feed back to update the classification.
3. A demonstration of the usefulness of such a model of social structure, pointing to phenomena that can be better understood as a result of deploying it, to relationships among phenomena it uncovers that were not previously obvious, and to observations of events or trends that can, given operative conditions, be predicted on the basis of it.

It is the purpose of this paper to provide such a defense. We concentrate here, at a fairly general level, on two major classes of social interaction – markets and science – but see no reason why a similar approach would not prove productive for the study of many other instances of institutionalized social arrangements, including city neighborhoods, firms, wiki collaborations, as well as (probably as instances of social systems lacking one or more of the typical characteristics of adaptive systems) bureaucracies, legislatures, legal systems, and money and banking arrangements.

2. General attributes of adaptive systems

In very general terms, an adaptive system is a network of interacting elements, immersed in a dynamic environment to which (at least some of) the elements are sensitive. The interactions between the elements which maintain the network are repetitive and localized, but not absolutely constrained, so that, while they give the network obvious structure, they can adjust, modifying the structure, as reactions to environmental changes percolate through the system.

The following is a general overview of the attributes of such adaptive systems, introducing concepts and vocabulary necessary for any discussion of the characteristics and capabilities of adaptive systems in the social domain:

1. The *active elements* of the system are *individuals* with a range of performance capability.
2. These individuals are sensitive to at least one *private benefit* which is scarce, individually valued, cumulative but not satiable, and to a significant extent performance-related.
3. The individuals perform by engaging in *interactions* which are capable of resulting, directly or indirectly, in acquisition of the private benefit dependent to some extent on judgments of performance determined by other participants in the interaction. These interactions are of definite and predictable form, and cost the participants some of their resources. Participants are constrained in their interactions by their resources available for interaction, their "nearness" to suitable partners, and the willingness and ability of potential partners to accept the interaction. Interactions that prove insufficiently beneficial (compared to other potential interactions) can be terminated relatively easily.
4. There is at least one *generalized benefit*, enjoyed potentially by all interacting individuals, which results from individual exposure to and participation in the system.
5. The *capability* of individual performance within the system is dependent on accumulated experience of interaction, accumulation of the private benefit, and exposure to the generalized benefit.
6. The system is open to its *environment*, and the environment's (sometimes novel) effects on individual elements can embody *incentives* for some of the elements to change their local interaction patterns. Interactions within the system can produce effects which change the environment, and effects from these changes can feed back into the system.
7. There are also *feedback* effects within the system, as the downstream effects of interactions can affect the incentives faced by upstream participants.
8. The interactions are sufficiently stylized and repetitive that they link the individuals into a *durable but mutable structure*. This structure can maintain its integrity while changing and even growing (incorporating new elements and interactions) within the environment.

9. At a level above the individual, *groupings* of similar interactions are recognizable and may have causal efficacy – their observation by other participants may influence the future actions of those observers.

10. The system's openness to its environment, in particular, the incentive effects on individuals to change their local interaction patterns to accommodate persistent environmental changes, make the system's structure sensitive to the characteristics of the environment which impinge on it. The major effects on the system's interaction structure will usually be quite local, but the structure-altering effects of follow-on interactions may propagate widely within the system. The changed structure is, in effect, the system's representation or *map* of the changed environment.

11. At any point in time, the activity within the system, comprised of ongoing interactions within the networked arrangement of individuals that comprise the system's map, is reflective of the environmental effects which are currently impinging on the system, and can be thought of as a *model* of the current state of the system's environment.

12. Particular classes of environmental effects will induce distinctive activation patterns in parts of the map, and these distinguished patterns or system states constitute the system's *knowledge* of the environmental effects. In other words, the system *classifies* the phenomena of its environment to which it is sensitive in terms of its internal states. Physically different phenomena which induce the same or similar states are recognized by the system as being the same. The categories formed are not inherent in the environment; they are "subjective" to the classifying system, and relational, comparable only between themselves.

13. It is essential to distinguish between the values, incentives, and purposes pertinent to the interactions of the participants from the apparent purposefulness or adaptability of the system as a whole and the values seemingly inherent in the system's reactions, the effects of which may feed back to affect the individual participants. The participants are each truly "led as if by an invisible hand to promote and end which is none of his intention".

All of the characteristics of adaptive systems described above are consistent with Hayek's (1952) description of the brain, with neurons as the active elements interacting physically via their axons, and the appropriate physical and chemical effects impinging on neurons as a result of their activity and location being the analogs of incentives and benefits. But, in generalizing from brains to adaptive systems in the social domain, let us be very clear that we are not saying that social systems *are* brains. We *do not* think that people and neurons are comparable in any other sense than that they can form mutable interaction patterns with each other. We most certainly *do not* think that social orders work like organisms with a brain, with one part set aside to specialize in "thinking" and exerting some considerable level of control over the rest – in fact, we can give good reasons in terms of our generalization why attempts to construct social arrangements on the model of an organism with a brain would have adaptive ability markedly inferior to particular arrangements that do not rely on central control. We *do* think that the basic processes of classification described by Hayek as operating in the brain, including particularly the formation of a mutable map of the brain's environment as experienced in the past and the ability of that map to support a model driven by current experience, have their counterparts in adaptive social orders, implemented differently, of course, but very similar in principle. Social orders are not brains but are brain-like, in certain very specific and circumscribed respects. The interactions between their components implement a classifying process on stimuli both external and internal to the system, and this generated "knowledge" of the environment has effects on component behavior and their structure of interactions that constitute adaptive reactions of the system to changes in its environment.

Yet, despite the obvious differences and qualifications, social orders have many characteristics in common with the neural order. The following are some distinctly brain-like aspects of social systems:

1. The systems have structural integrity while at the same time exhibiting a mutability in response to environmental experience. The pathways of interaction within the system are affected by environmental change in characteristic ways – fundamental, long-used interaction patterns, such as the institutions of property and exchange in markets and publication and citation in science, are very resistant to change, whereas the more peripheral, more recently acquired, less pervasive ones are much more easily modified or discarded. The structural rearrangements are path dependent and irreversible, with any change altering the context in which the system reacts to subsequent change.

2. One can observe an expansion of number of elements involved in pathways associated with a particular stimulus as that stimulus becomes more frequent or more pressing, and a concomitant reduction in numbers of elements associated with relatively less pressing stimuli. This is the analog of "competitive plasticity" observed in brains. The effect shows up most noticeably as "creative destruction" in markets and "paradigm changes" in science. When a particular type of stimulus has not been experienced for

some time, the pathways associated with reactions to it will tend to atrophy, with any permanent resources involved being devoted to other interactions. An obvious example of such "forgetting" is seen in the relatively small part of a modern economy devoted to horse technology.

3. The systems exhibit functional mutability – the foreclosing of interaction pathways (by whatever means, by natural events or by intentional intervention) does not, in general, permanently cut off the response to the stimuli which activated these pathways. Instead, the system tends to work around the impediment, often co-opting of neighboring pathways previously used for servicing other inputs. Two examples of this effect are given below in the sections on market intervention and economic history.

4. The systems exhibit anticipation in the sense that their current structure predisposes them to react to future events in ways dictated by their experience of past events. Examples of such anticipatory capability are seen in the futures prices in markets and the predictive aspects of scientific theories. Its conservative side (the assumption that the future will be much like the past) shows up, for example, in the tendency of science to focus on phenomena consistent with an accepted paradigm and to ignore conflicting data. An interesting example of this characteristic of seeing only what is expected is described in Pickering (1984: 301-302) where experimental data reported by some *atomic* physicists inconsistent with electroweak gauge theory was dismissed by *particle* physicists without any specific flaws in experimental technique ever being suggested.

5. The systems exhibit reactions analogous to addiction – with cash or credit (in markets) or funding (in science) or internal security (in firms) as the analog of neurotransmitters in brains. In markets and science, the ensuing "hangover" of recession and readjustment are predictable features of the recovery process.

6. The classifications performed by these systems are relational, and what are being related in the final analysis are internal states of activation. The classes are the focal interaction patterns that the current configuration of the system's map supports. These particular patterns of activation, learned from past experience, are the "attractors" which form the basis of the system's classifications of future inputs. To give an illustration, if one observes the interaction patterns of Austrian economists – in papers, conference speeches, and conversations – one sees that the notion of capital as a structured entity is deeply embedded in these interactions. One could say that the scientific subsystem of Austrian economists knows that capital is a structure (or is fruitfully modeled as such). Confrontation with

empirical phenomena involving capital generates interactions in which structured capital is a feature. Interactions characterized by structured capital tend to get follow-on notice; those with an unstructured, "capital stock" characteristic tend to die out rather quickly. Structured capital is an "attractor" in the interaction space of Austrian economists.

7. There are areas of interaction that are within the boundaries of the system but which are confined to a local context and different from the interactions characteristic of the system itself, analogous to transmissions within the cerebellum or the basal ganglia within brains. Examples include production activities in firms within markets, research activities in labs within science, and interactions in families within markets.

8. How a system reacts to a particular stimulus depends not only on the stimulus but also on the current state of the system. For example, in markets, there is no guarantee that small minimum wage increases will cause unemployment increases – one needs to look at the available modes of adjustment in the current circumstances. Simple supply-and-demand analyses abstract from this context and assume no constraints on adjustment.

Suggestions that economies and other social arrangements can be understood as examples of "complex adaptive systems" have been made many times before in complexity theory – see, for example, Kauffman (1993: 395-402) – and sociology – see, for example Buckley (1998). It is also implicit in Hayek – see, for example, Hayek (1967: 66-81), where he asserts that "there is no reason why a polycentric order in which each element is guided only by rules and receives no orders from a center should not be capable of bringing about as complex and apparently as 'purposive' an adaptation to circumstances as could be produced [in a more hierarchically organized system]". Kauffman (1993: 173-235; 1995: 71-92), discussing "the twin sources of order", makes the important point that, in biological systems, natural selection is not the only source of order, for the tendency in certain systems to self-organization is, in a sense, "order for free". By this he means that the formation of that sort of order is 'spontaneous', a designation that will resonate with those familiar with the work of Hayek – see, for example, Hayek (1973) and Boehm (1994). But systems involving some form of self-organization are members of a very large and diverse set, and so the commonalities are likely to be of such a general nature as to provide very little assistance in understanding particular social systems. So, while endorsing Kauffman's insight, we prefer to considerably narrow the field so that more concrete things can be said about the structure of the systems of interest. Our purpose here is to push beyond such programmatic statements by injecting more explicitness into the idea.

Note that adaptive systems are not simply ensembles that adapt as a result of their environment's selection pressure in promoting the existence and reproduction of some components and discouraging others, like a species or an immune system (although their components may, indeed, be subject to such selection). There is a categorical difference between groups that react (successfully) to environmental phenomena through the additive actions of individuals and those whose reactions are a result of (and induce) learning at the group level. The individual behaviors that promote successful reaction in the first type of group may evolve via selection; the learning characteristic of the second type comes about through changes to the mutable internal interactional structures within the group. While most research on group behavior (flocking, path formation, etc.) deals with groups of the first type, it is the second type – which includes systems as diverse as ant or termite colonies and the brains of higher animals and, we claim, various systems of human social interaction, including markets and science – which interests us here.[3]

3. Markets and science as adaptive systems

The term "science" is used here to refer to the complex of people and institutions that make up the knowledge-generating activities of a scientific community rather than, as might be more common, the knowledge itself that is generated within that social system. Similarly, "market" refers to the complex of people and institutions that make up a community of buyers and sellers in a money economy. Not included under "market", however, is activity that takes place within firms or families, and not included under "science" is the activity of academic teaching or the interpersonal interactions occurring within laboratory teams. In the context of these definitions, science and market are distinct instances of adaptive systems in the social domain.

The terms "market" and "science" as used here do not refer to any particular market economy or arena of scientific endeavor. That level of specificity ("the U.S. economy", for example, or "the physics community") would be necessary in applications of the basic theory but would be unhelpful to the task of setting out the basic concepts. The fact that applications might be dealing with more than one market or more than one scientific domain does, however, raise the analytical possibility that, for example, part of the environment of the market of interest is another market – a division into system and environment that could be effective if the amount of interaction between the two markets was relatively small compared to the activity in the market of interest.

3.1 Markets

A market, viewed at a very high level of abstraction, is a complex of people interacting via goods exchanges. A market system's structure – its Hayekian map – involves the following elements:

1. The institutional framework of property, contract, and money – the fundamental and long-lasting institutions generally adhered to by all participants without which market activity on a large scale would be infeasible. Describing these institutions as a "framework" is not meant to imply that they are immune to change – indeed, they could not have and did not arrive on the scene in their current form; they are the result of a long process of development. As Hayek (1945: 528) puts it: "We have developed these practices and institutions by building upon habits and institutions which have proved successful in their own sphere and which have in turn become the foundation of the civilization we have built up."
2. The personal habits and routines that market participants have learned to rely on to implement their plans. People's activities tend to follow generally repetitive patterns of interaction, and although deviations may occur and action details may vary, the underlying routines are not, in ordinary circumstances, subject to dramatic change.
3. The market participants themselves, the active components of the system, whose exchange transactions involving goods and services constitute the impulses that animate the system, whose tastes and preferences (the satisfaction of which constitutes the private benefit available through market interaction) can be changed as a direct result of market experience, and whose accumulations of wealth and "brand name" affect not only their capabilities for interaction but their tastes and preferences as well.

The current structure of a market – the complex of market participants connected to specific other participants via their normal constellation of exchange interactions – is the result of the system's past experience of its environment. An obviously important part of this environment consists of the natural resources available to market participants. The stimulus of the reduced quantity available of a resource is "felt" by the system, with the "feeling" being a rise in relative price of the resource and, to varying degrees, of its various derivatives. Many different physical events (droughts, floods, bad weather, wars, for example) can have the same supply-reducing effect, and the system does not discriminate between them. The same physical event can also be sensed as different in different circumstances, e.g., high winds can increase the prices of fruit products but decrease the price of electricity. Since goods derived from the original resources are also classified, the overall spectrum of market prices represents a very fine-tuned classification of the original inputs and their relationships to each other. In addition, the environment also includes effects from other social arrangements, for example, those that deal with contract enforcement and those that control the supply of money.

The market's model of its current environment is the ongoing flow of transactions (characterized by transfers of goods and observable exchange prices) between the market participants. These transactions are induced by stimuli from environmental inputs and conditioned by the preferences and creativity of the market participants themselves. They follow transactional paths constrained by the current structure of the map, and they result, indirectly, in a classification of the various stimuli currently impinging on the market system, a classification visible to the participants in the array of market goods and their market prices.

This classification is an emergent phenomenon. An individual may intend to buy or sell a particular good, but no individual plans the overall configuration of marketable goods and services, related to each other (as an emergent result of market activity) as inputs and outputs and as complements and substitutes of varying degrees. An individual may deliberately set a desired price, and an exchange may be consummated at some particular exchange price, but no individual plans the emergence of the spectrum of market prices that relate different goods and reflect overall appraisals of desirability and scarcity. Yet, the market system could not survive as a coherent system unless these market goods and market prices were to some extent an operational reflection of actual resources, scarcities, needs, preferences, and the concomitant constraints imposed by other social systems. And because this spectrum of prices is a sensitive reflection of environmental conditions and is, in principle, visible to all participants, it constitutes a generalized benefit for all participants in the system.

Market participants can observe not only their own transactions but those around them, and they can read reports of transactions others have observed (such as quotes of stock prices). They are also recipients of advertising from vendors apprising them of potential transactions. Based on their appraisals of this information, they can modify their own transaction repertoire, perhaps, for example, initiating transactions for a good that they see a lot of other people buying – transactions which, from the observer's point of view, represent potentially appropriable gains. In this way, novel local stimuli can have systemic effects, altering the system's map. Most such alterations will be at the level of changes in individual preferences and minor amendments to personal routines, although even relatively short exposure to stimuli that induce transactions in which there are large and obvious (at least to an entrepreneurial observer) unappropriated gains can result in more substantive change. And extended, unresolved exposure to such stimuli can result in changes to the more stable areas of the map – as, for example, during the 1100s, when contractual conventions supporting exchange were expanded to encompass negotiable credit instruments (see Benson 1989), a development with cascade effects, leading in turn to many changes in commercial activity including, eventually, the emergence of formalized futures markets. In any case, the array of marketable goods and services, the quantities brought to market, and their market prices will adjust as a result of even rather minor changes in the map and therefore will form a highly detailed and sensitive classification of the environmental influences experienced by the system.

The market as a system can exhibit anticipation of environmental phenomena and active intervention to cope with the expected effects. Consider, for example, an exogenous environmental change in the form of a removal of some prohibitions on offshore oil drilling. Although this has no effect on the current supply of oil, and no direct effect on the current prices of oil and oil-derived products, it will trigger intertemporal adjustments as individual traders speculate on future supplies and prices. And, as Coats and Pecquet (2008) point out, reductions in futures prices will feed back to affect the individual decisions of companies with currently extractable deposits as to their rates of extraction. The incentive will be to increase rates, and so increase current supply, tending to lower current prices. The systemic result of this interplay of individual transactions is, first, a spectrum of market prices that is reflective not only of the current state of the system's environment but also of the expected future state, and, second, action taken to mitigate the adverse consequences of the expected future state.

3.2 Science

At a similarly very high level of abstraction, science is a complex of people interacting via the institutions of scientific publication and citation. A science system's structure – its Hayekian map – involves the following elements:

1. The institutional framework of publication and citation – the fundamental and long-lasting institutions without which science on a large scale would be infeasible. These institutions appear to have taken their current form as recently as the late 1600s, with the advent of the Royal Society and its journal, the *Philosophical Transactions* – see Merton (1973: 191-203, 460-496), Hull (1988: 323-324), and McQuade and Butos (2003).

2. The personal habits and routines that scientists have learned to rely on to implement their plans. These generally repetitive patterns of interaction include their organization into schools and groups and their patterns in selecting their usual outlets for communication and publication.
3. The scientists themselves, the active components of the system, whose transactions involving publication of articles, use of information in published articles, and citation of information used, constitute the impulses that animate the system, whose tastes and preferences (the satisfaction of which constitutes the private benefit available through scientific interaction) can be changed as a direct result of their experience as scientific researchers, and whose accumulations of reputation and credibility affect not only their capabilities for interaction but their tastes and preferences as well.

The current structure of a science – the complex of scientists connected to specific other participants via their normal constellation of publication and citation interactions – is the result of the system's past experience of its environment. An obviously important part of this environment is the physical world, with experimental scientists (and their equipment) as sense organs. But the environment also includes effects from other social arrangements, most importantly those that fund scientific activity.

Science's model is the ongoing flow of transactions (characterized by publication in various forums and citation when invoking the work of others) between scientists. These transactions are induced by stimuli from environmental conditions (including experiments and observations) conditioned by the preferences and creativity of the scientists themselves. They follow transactional paths constrained by the current structure of the map, and they result, indirectly, in a classification of the various stimuli currently impinging on the system. This classification, experienced by the participants as what would be characterized by Kuhn (1962) as a "paradigm" (in the sense of the current state of science or of a branch of science), and by Kitcher (1993: 87-89) as "consensus practice" can, in part, be described in propositional form as the theoretical and taxonomic corpus of established scientific knowledge.

A science's classification, like that of a market, is an emergent phenomenon. An individual may intend to develop and expound on a particular theory, but no individual determines how, and in what form, the insights of this theory are incorporated into the current body of scientific knowledge. An individual may deliberately seek reputation, but no individual is in control of the emergence of the overall assessment of his reputation which reflects appraisals of the value and usefulness to others of his contributions. Yet, science could not survive as a coherent system unless the body of scientific knowledge was to some extent a useful classification of natural phenomena, and its operation would certainly be seriously hampered if the general assessment of reputation persistently ignored important contributions. And because this corpus of scientific knowledge is a useful base on which to build further individual research and is, in principle, visible to all participants, it constitutes a generalized benefit for all participants in the system.

Scientists are concerned not only with their own work and that of their colleagues and competitors but with the work of those in related fields – they probably read (and absorb information from) many more articles than they cite. (This is not to imply that scientists chronically avoid citation; it is simply the case that, since citation only occurs in published articles employing the published results of others, there is a lot of scope for more subtle influences to be absorbed from articles read.) They can also be quite attentive to who has won prizes, which areas of research are 'cutting edge', and where the most lucrative grants are to be had. Based on their appraisals of this information, they can modify their own transaction repertoire, perhaps, for example, turning attention to a phenomenon that they see as attractive to investigate due to lack of competition or availability of funding – transactions which, from the particular scientist's point of view, represent potentially appropriable gains. In this way, novel local stimuli (unexpected observations, for example, or new sources of funding) can have systemic effects, altering the system's map. As is the case for markets, most such alterations will be at the level of changes in individual preferences and in personal routines, although even relatively short exposure to stimuli that induce transactions in which there are large and obvious (at least to an entrepreneurial observer) unappropriated gains can result in more substantive change. And extended, unresolved exposure to such stimuli can result in changes to the more stable areas of the map – as seen, for example, in the emergence of working paper circulation networks and (more recently) the trend toward internet posting of articles, both of which are responses to the perceived costs of the lead times experienced in the conventional publishing process. In any case, the current body of scientific knowledge will adjust in response to both major and minor changes in the map, and (perhaps more slowly and less perceptibly) the complex of reputational assessments will change also.

Science as a system can also, like a market, display anticipation. In the absence of any experience to the contrary, this anticipation tends to take the form of a systemic expectation that the future will be much like the past. This tendency can be seen in the operation of research groups or "schools" and their training of graduate and postdoctoral students in the techniques, presumptions, and core ideas of their field. The effect of school identification and training is to condition the scientists involved to be sensitive and receptive to certain inputs from the environment, to be selective with regard to the appreciation of contributions of other scientists, and, generally, to view the stimuli they encounter through the filter of their school's presumptions. In this way, the organization of scientific schools and training is an embodiment of systemic expectations about the character of the environment, based on previous experience. This does not mean that surprises cannot occur and even lead to alterations in the map in the form of reorganizations of school affiliation or changes in core ideas. But, as in the sensory domain where unexpected input can easily be completely ignored, science can, for long periods, proceed in ignorance of phenomena (continental drift, for example) that have been detected but which have been filtered out as contrary to expectations, and it is only with the experience of repeated stimuli that adaptation finally occurs.

4. Example: interventions in markets

The topic of intervention in markets is a well-rehearsed one in economics; for a definitive book-length treatment in the Austrian tradition, the reader is referred to Ikeda (1997). Adaptive systems theory is quite compatible with Ikeda's emphasis on the involvement of system-wide endogenous processes in the analysis of intervention as well as his highlighting of its downstream, unforeseen consequences. In discussing the effects of intervention, Ikeda highlights, in addition to the incentive effects which are the usual subject of economic analysis, what he calls "discovery effects" – effects which result from redirection of entrepreneurial activity away from some areas ("stifled discovery") and toward others ("superfluous discovery").[4]

What adaptive systems theory adds to Ikeda's analysis is a more comprehensive, structural picture of the system being perturbed by interventions. In general, interventions represent either a change in the environment of the system or an alteration to the operation of the system by foreclosing or redirecting particular patterns of interaction. Either way, an adaptive system will adapt, changing its interaction patterns (which include "entrepreneurial activity") to recognize the presence of the external change or to work around the internal alteration. This adaptation will result in a new spectrum of market prices, accompanied by changes in the usages of the resources employed in market interactions. Whether the new configuration is better or worse than the old based on some criterion of individual welfare, adaptive systems theory does not say. What it can do is to describe the consequences of the change for the stability and adaptability of the system in facing future potential environmental changes.

The following examples should give the flavor of the perspective offered by adaptive systems theory in describing the nature and effect of particular types of market intervention:

1. This is an example featuring an internal intervention resulting in unintended local consequences. One possible result of a constraint on the ability of participation in a particular interaction will be a switch to participation in some available alternative, giving an interaction configuration not likely to be predictable by the intervener. Interventions that foreclose the use of some particular institution will tend to provoke a move to the adoption of a "nearby" institution (one known to the participants in a different context) for the purposes served by the foreclosed one. Examples of this phenomenon are not hard to find. The adoption of the institution of "shoe money" (traditionally, a payment by *prospective* tenants to pay a middleman for help in negotiating a rental contract) in arrangements with *existing* tenants in the aftermath of Hong Kong's imposition of rent control in the 1920s, described by Cheung (1975), is a clear example. More generally, the phenomenon of the market's map adjusting to the shutting down of existing transactional paths has been recognized elsewhere as "intervention breeding more intervention" – see, for example, Mises (1949: 762-764) and Ikeda (1997). There *will* be a workaround subsequent to intervention (since classifying systems can do nothing else but classify) and, if (as often happens) the ensuing state of affairs is unsatisfactory to the intervening party and the response is further intervention, the result is simply further injury to the adaptive capability of the system.

2. This is an example featuring an external intervention resulting in unintended global consequences. Unless the intervener knows exactly how a change in the system's environment affects the global classification (which is most infeasible in a complex system), the global classification will change in unanticipated ways. These changes will feed back into local interactions, producing further unanticipated changes at the local level. For an example of such generation of systemic uncertainty, consider the phenomenon of markets with "big players" (as described by Koppl 2002), i.e., functionaries, such as central bankers or finance ministers with discretionary power to affect market conditions, who are outside the market system in that they are, in these roles, immune to the discipline of profit and loss. The presence of a big player introduces a prominent element into the environment of the market system, and so the resulting classification is an adaptation not only to the usual environmental conditions but also to the big player's actions. Especially in cases where the big player's activities affect the supply of money and credit (since money is a component in every transaction and many

transactions, involving both capital and consumption goods, are dependent on credit) the market's adaptation is significantly affected by the actions of the big player. For a normal market participant, it is at least as important to predict the future behavior of the big player as it is to predict the "underlying fundamentals". If big players engage in discretionary policy, market participants are not likely to be able to adjust their transactional behavior quickly enough to the realities of the new (and ever-changing) environment and so the market classifications will be only poor representations of the actual circumstances. Since, in the big player market, the market classifications are much less useful for augmenting individual knowledge of the circumstances beyond those affected by the big player, the opportunities for market participants to adapt their individual knowledge to these circumstances is impaired. Their obvious incentive is to adapt their knowledge to the behavior of the big player, but this can be a difficult and error-prone process in a regime of discretionary policy. In looking for clues here, market participants are likely to form expectations heavily based on what others think the big player is going to do, thereby providing a rationale for "herding" behaviors.

5. Example: science funding

The phenomenon of science funding provides a case study of the reaction of a particular adaptive system to strong influences from outside the system. In order to apply the theory of adaptive classifying systems to science funding, one needs first to get a clear idea of what constitutes the process of "ideal science" – science floating above funding considerations. (This is analogous to the way one conceives of "ideal markets" – markets floating above security and enforcement considerations.) Then, introducing real systems of funding, one investigates what effects these can have on the pristine process of scientific publication, citation, and reputation formation.

To take two polar cases, science can be embedded in market – with private funding – or it can be embedded in government – with public funding. Both embeddings are nonneutral with respect to the operation of science, and the puzzle becomes one of figuring out the different effects each is likely to have on the structure of the overall system – effects on the incentives and the feedback mechanisms which drive the knowledge-generating process, and on the ability of the participants to be attracted to and to remain as contributors to the process. The questions that follow naturally from this set-up are interestingly different from those normally considered in the economics of science.

Here are some examples of the sorts of questions that suggest themselves:

- Is the nonneutrality simply one of pushing different lines of research within the parameters of science itself, or is it one of affecting the structure of the knowledge-generating process?
- And, if the latter, are the new incentives (to choose potentially profitable areas of research or to choose politically favored, pressured, or mandated ones) detrimental to the reliability of the knowledge generated?
- Are reputational feedback effects changed so that the importance of peer acceptance is diminished and contributions do not go through the reinterpretation process to the normal extent?
- Are the criteria for peer acceptance changed? Is scholarly reputation no longer the only form of reputation sought?
- Does the process of obtaining funding itself have psychological effects on the individual scientists that might affect their participation in the science process?
- Are the characteristics of the funding process such as to encourage stable growth in the body of science and scientists or to be more inclined to induce episodes of boom and bust?

We are interested to know how the scientific order functions and what happens to it under an institutional framework characterized by significant government funding of R&D – a regime characterized by a very small number of very large sources of public funds, typical of the current situation in the U.S.[5] As can be seen from the above questions, this is a non-normative approach to the study of the scientific order and geared to its functional, adaptive, and emergent attributes as a knowledge-generating social order, in much the same general sense that one might expect economists to study of the market order. But it is important to emphasize that the principal institutions of the scientific order are not market institutions – they do not generate market prices, and they do not directly produce exchangeable commodities. Analyzing science as if it were a market process obfuscates the operation of the relevant institutions of science and promotes the use of evaluative criteria possibly of use in the theory of catallactic systems but unsuited for the positive analysis of science.

In Butos and McQuade (2006) we argue that a funding arrangement with large, concentrated, and politicized sources of funding power has the potential to alter the adaptive properties of the scientific order – in the sense that it tends to erode the corrective action of the basic institutions of science by introducing new incentives, adding political considerations to the reputation-assessing mechanisms, and generating short-term research enthusiasms in a way very reminiscent of boom and bust in markets. Such effects are much less likely and considerably more constrained if the funding is decentralized and free of political oversight. More generally, three types of effect can be identified which principally concern how monopolistic funding affects science:

1. The *direction* of science in terms of favored R&D projects and research questions.

2. The *destabilization* of science in terms of the effects on inputs used in science induced by the pattern of funding. In some cases, this can take the form of a boom and bust sequence reminiscent of the analogous phenomenon in markets. The analog of capital processes in science may involve investments that come to fruition in the publication of papers. The building of research labs and school facilities would be relatively far removed from fruition; the acceptance of students into graduate programs would be relatively closer. Availability of funding would drive a boom in any particular subdiscipline, and the crunch would come as resource restraints became pressing or as the attention of funding agencies moved elsewhere. The crisis would be marked by significant increases in lobbying and funding-seeking activity. The downturn would involve unemployment (within the subdiscipline) and facility underutilization.

3. The *distortion* of the procedures of science in the sense that concentrated funding power may promote a knowledge-generating and certification process not consistent with the evolved mechanisms of science – in particular, the publication-citation-reputation interactions described above (henceforth, "PCR") – that confer scientific legitimacy at the system level to the knowledge produced by scientists.

It is important to note that the first two effects do not imply a systematic bias in the scientific quality of the research or its findings whether the funding source is public or private. All funding sources necessarily affect the kind of science undertaken, and just because government funds the great majority of basic science in the U.S., there is no necessary implication that any directional reorientations of science, even if these involve seemingly perverse changes in allocation, produce inferior knowledge relative to what would have been generated had funding been directed in some other way or provided by other sources. Nonetheless, it is likely to be different. But as long as the basic procedure for knowledge certification – the PCR process – remains uncompromised, quality is not an issue. In short, funding, from whatever source, will affect the direction of research and the utilization of resources and hence the content of scientific knowledge, but this in itself does not mean that the result will not be good science.

This is not the case, however, when distortive effects are introduced by changes to or replacement of the PCR process. Whatever directional and stability effects may arise from science funding, there is now a separate compounding issue if the institutional structures of the scientific order that legitimate theories and empirical findings are compromised, bypassed, or replaced. It is one thing, for example, if a government agency only redirects scientific inquiry by changing funding from Project X to Project Y (from space science to 5th generation computing, to cite a specific instance); it is quite another if the government enforces, for example, a decree that certain directions of research are to be funded and scientific results are to be accepted or rejected based on their compatibility with certain ideological propositions, or, more subtly, if a scientist's political connections convey scientific credibility over and above that earned from scientific research.

From an empirical standpoint, funding mechanisms having the capacity to distort the evolved procedures of science are rarely so explicit or severe that the destruction of science may be unambiguously attributed to them, although the era of Stalinist biology would certainly be an example. Instead, distortions of the PCR process, especially when they are unintended, tend to be subtle and difficult to observe or document. And even though science is ordinarily constrained by procedures comprising PCR, there is no reason to assume that scientists do not have incentives to chisel, cut corners, or engage in unacceptable practices in the hope that additional reputational gains are conferred on them. Yet, distortions of PCR and disreputable behavior of individual scientists become threatening to the scientific order only if such deviations and practices are systemic or become embedded within the conventions and institutions of science. In the real world, the juxtaposition of possibly all three effects presents an empirical challenge in identifying and documenting the separate independent effects and their linkages to each other. Finally, it is necessary to acknowledge that each of these effects reflect both intended and unintended byproducts of science funding.

A serious problem with monopolistic public funding is that it leads naturally to calls for "oversight" by politicians of the funded science. In the recent Baltimore Case (described in detail in Kevles 1998), the oversight prerogative was exercised in ways that moved beyond an investigation into misconduct and into the area of actually judging scientific disputes. The paper that thrust the scientists Baltimore and Imanishi-Kari into the national limelight had survived the journal's editor, peer reviewers, and two university committees of inquiry; nonetheless, determining whether it constituted good science emerged as relevant to the Dingell Congressional oversight committee. The Committee equated scientific error with fraud, thereby inferring that Baltimore and Imanishi-Kari acted in bad faith and did so with fraudulent intent. In terms of the PCR process, the intervention of a government agency in adjudicating scientific disputes would be seen as an imposed change in the "publication" component, though having implications for the citation and recognition components, as well. Government funding of science is not separable from government regulating science; the Baltimore Case reveals that the good (and not so good) intentions to uncover fraudulent scientists poses a strong risk to science of also inducing distorting elements in the scientific order.

But rather than any particular result, however, what we would like to emphasize is that this approach to the economics of science completely changes the set of questions one is driven to ask. Science is not taken to be a somewhat peculiar and dysfunctional type of market – it is an adaptive system in its own right with its own particular complement of interactions and incentives. Prescriptive idealizations such as welfare optima, which are problematic even for markets, have no valid application to science. No particular method of funding is to be taken for granted – different funding regimes will have different effects on the system, and these effects can be investigated both theoretically and empirically. The emphasis is less on attempting to quantify outcomes in the sense of research "production" (as if science were a commercial firm) and more on attempting to characterize outcomes (probably completely unintended ones) in the sense of structural effects on the processes of science itself.

6. Example: economic history

This application of adaptive systems theory to economic history invokes an analogy to what happens in the brain when some neural pathways involved in the classification of particular stimuli are closed off by injury. The brain is not passive, and in continuing its classification process it forms new connections which work around the damage, connections which tend to involve nearby parts of the map that were previously active in the classification of other stimuli. A similar effect should be evident in social systems, and so our hypothesis is that, in understanding more deeply the reactions that, historically, have occurred subsequent to changes that foreclose some previously operational transactional pathways, one should look closely at the contemporary institutional surroundings.

The particular historical episode of this example involves the period in 14th century Florence subsequent to the Ciompi Rebellion of 1378 which involved members of the woolen industry guild demanding protection from competition from the emerging English woolen manufacturing industry. The political reaction to the rebellion was to suppress it pretty effectively but, in the process, the existing guild-based system of commercial relations was thoroughly – and, as it turned out, permanently – disrupted. But, in a very short span of time, the guild-oriented commercial order was largely replaced by new cross-industry organizational forms which had the effect of revolutionizing Florentine commerce through the widespread use of credit and fostering the development of new industries, particularly merchant banking and silk manufacturing and retailing.

Accounts of this period by sociologists and historians emphasize what they refer to as the "social embeddedness" of the new commercial interactions. By this, they mean that the new arrangements of business partnership and commercial credit were inextricably tied up with political, family, and friendship connections. To some extent this is quite accurate, but it is far from being the whole story. For example, a significant number of the partnership arrangements did not involve family or even people living close by (and were even established between people supporting different political factions), although they did consistently involve people of the same social class. In contrast, credit relationships frequently were established between people of different classes, but tended to involve inhabitants of the same neighborhood.

We suggest, following Padgett (2001) and Padgett and McLean (2002), that the new arrangements which emerged in the aftermath of the destruction of the guild-based business environment resulted from the co-opting of existing institutional forms – business partnerships taking on the characteristics of the institution of upper-class marriage, and commercial credit relationships taking after the institution of "making" friends or clients (in the Florentine literal sense of "making" – developing and sustaining – which apparently did not necessarily imply affection but which did involve repeated transactions over an extended period of time).

In the case of business partnerships, this is consistent with their emergence among upper-class businessmen, a development which was subsequently copied by businessmen in other social strata but which still retained the strong within-class characteristic. Documents from the time show considerable overlap in these people's minds between the norms of marriage and those of business partnership, to the extent that it was not uncommon for one partner to take care of the other partner's children in his will. In the case of commercial credit, the combination of within-neighborhood and across-class characteristics reflects the sense of obligation inherent in the Florentine concept of "neighbor" which underlay the friend-making institution.

In summary, this episode seems to be a clear example of commercial behavior co-opting institutions in a "nearby" area of interaction in the aftermath of a disruption that shut off the operation of existing commercial institutions – a response in both kind and particulars that one would expect in an adaptive classifying system.

7. Example: epistemological considerations

One of the characteristics of modern (postpositivist) philosophy of science is the adoption of an appreciation for the relevance of actual scientific concepts and findings, especially about human beings, and the history of actual scientific discoveries and the methods by which they were produced. A second feature is the appreciation for the relevance of social factors in the generation of knowledge, a concern that lends itself to a naturalistic approach. Entangled with both of these developments is the question of whether or not a normative stance, the traditional stance of epistemology, is still feasible. And further stirring the postpositivist pot is question of whether, in the absence of any *a priori* foundation for a normative stance, the way is left open for a relativism in which any sort of belief formation is warranted.

For the purpose of bringing adaptive systems theory to bear on this wide field of philosophical activity, it is useful to categorize the varieties of postpositivist naturalism into individually-oriented and socially-oriented, and to divide each of these further according to whether or not their thrust is normative.

	Individualist	**Social**
Nonnormative	Quine, Churchland	Barnes, Bloor
Normative	Laudan	Goldman

Here is some elaboration on each of the four possibilities:

- The original nonnormative individualist naturalism is that of Quine (1969), who suggested that traditional concerns about justification of beliefs confused two separate issues: explanation and legitimization. The latter is a matter for cultural studies; the former, focusing on a causal understanding of an individual's belief formation, is the province of psychology. This either eliminates any role for traditional epistemology, or reduces and reorients it as a branch of psychology. It is not a popular stand among postpositivist philosophers, although it has had some downstream influence. From the point of view of Hayek scholars, a particularly interesting variant of nonnormative naturalism is that of Churchland (1995), based on a theory of mind as patterns of activation in a recurrent Hebbian neural network which, although apparently pursued in complete ignorance of *The Sensory Order*, has much in common with Hayek's work – while going beyond it in adding considerable neurobiological detail, experimental grounding, and philosophical sophistication.

- Far more mainstream is normative individualist naturalism, of which the work of Laudan (1990) is an influential example. Here, the findings of science and the history of science are mined for clues as to what constitutes good reasons for belief and good practice. Instead of seeking to formulate *a priori* standards against which to evaluate scientific work, scientific activity is taken as paradigmatic of rational behavior. The aim is to pick out, from what scientists have actually done, the strategies which have made possible the emergence of science as a successful enterprise. The naturalistic fallacy – the construction of "ought" from "is" – is skirted by keeping any proposed norms strictly hypothetical.

- A prominent example of nonnormative social naturalism is the philosophical stance of Barnes (1974), Bloor (1976), and their followers in the Edinburgh School of sociology. Their claim is that they provide an exclusively descriptive account of scientific knowledge while adhering to their four general tenets: causality (the intent to uncover the causal origins of beliefs), impartiality (the commitment not to privilege particular belief-generating systems), symmetry (the explanations of beliefs considered false invoking the same causal mechanisms as those of true beliefs), and reflexivity (the acknowledgement that the beliefs of the investigator are subject to the same analysis as those of the investigated). Knowledge is socially constituted in that norms of rationality emerge in specific social contexts and subject to the effect of interests from outside of science as well as inside, and knowledge is taken to be whatever beliefs are collectively endorsed in a particular social context. This relativist conception of knowledge would seem on its face to have ignored the rather important matter of the constraints imposed by the physical world itself, but Bloor (1976: 7) does acknowledge, if somewhat vaguely, that "naturally there will be other types of causes apart from social ones which will cooperate in bringing about belief".

- Probably the best-known normative social naturalism is Goldman's (1999) "social epistemics". Goldman proposes to associate a "veritistic value" to states of belief, where belief in a truth rates highest on the scale. He is a strong realist about truth; something is true if it conforms to what is observed to be so. Practices (including norms, rules, and forms of interaction) in a scientific (or any other) community can have positive instrumental veritistic value if they tend to increase the mean veritistic value across the members of the community, and this measure allows for a comparative assessment of existing and proposed practices. The evaluative standards proposed – reliability, power, fecundity, speed, and efficiency, according to Goldman (1987: 128-129) – are universal ones, not relative to any particular community, although their respective weights may differ in different communities. And Goldman does not imply that some particular practices are always to be used, for a range of diverse practices operative within a community may turn out to score higher on the veritistic scale than uniformity of practice.

Since adaptive systems theory deals with individuals embedded in a social system, studied in a way that has more in common with anthropology and sociology than epistemology as traditionally understood, it is no surprise that it would be closer in spirit to social naturalism than to individualist naturalism, and more in sympathy with nonnormative naturalism than with normative naturalism. However, there is no argument with nonnormative individualist naturalism as far as it goes, for understanding the behavior and reactions of individuals engaged in knowledge-generating activity is clearly important for understanding the system as a whole. This is particularly true with respect to Churchland's work, since this is, in effect, carrying further Hayek's work on which the generalization to social systems pursued here is based. Similarly for its normative counterpart – investigations of individual scientific practice, past and present, are useful inputs to a theory of science even if the motive for pursuing them is to unearth some normative pronouncements.

Nonnormative social naturalism is the species of postpositivism closest to adaptive systems theory. But there are considerable differences between the two both in emphasis and in substance. The Strong Programme has been taken to task, for example by Goldman (1987: 112), for its concentration on factors influencing knowledge generation external to science, particularly ideological and political influences, and on its neglect of the world itself as a major factor. But from the perspective of adaptive systems theory these objections miss an essential point – that it is not so much the interests and motivations of the individual scientists that determines what emerges as knowledge but the fact that these interests and motivations have to be played out in the context of interactions which incorporate very specific feedback effects on those individuals. They are the "invisible hand" of science. And these interactions, particularly the fundamental ones surrounding publication and citation, do not simply influence scientists' decisions and choices or constrain them to do the right thing, but are an integral part of the process by which individual contributions are transmuted into scientific knowledge. The adaptive system of science is built from *both* the individual scientists and their knowledge-generating interactions, and these in combination determine which external phenomena the system as a whole is sensitive to and to which it adapts. With an emphasis on usefulness based on empirical success built in to the basic interactions, the physical world is certainly top of the list, but social factors, including, for example, funding, can have their effect as well. Where adaptive systems theory radically departs not only from Bloor and Barnes but also from all of the above naturalisms is in its insistence on a categorical difference between individual knowledge and scientific knowledge, between knowledge instantiated in the internal structure of a brain and knowledge instantiated in the internal structure of the social system of science.

What all of the above forms of postpositivism have in common (to widely varying extents and levels of emphasis) are the philosophical claims of "underdetermination" (of theory by data), "incommensurability" (of prior and current or of competing scientific theories), and "theory-ladenness" (of data and methods). But these claims are always given an individualistic interpretation, whether in the context of theory choice (by individual scientists) or the analysis of observation (again, by individual scientists). However, from the perspective of adaptive systems theory, they are also systemic characteristics of scientific knowledge – they are a consequence of the way in which scientific knowledge is generated and the form in which it is instantiated in the structure of the social system. All of them arise from the way in which new inputs, incompatible with the system's current "map", to the extent that they are not simply ignored, result first in local adjustments, only leading to deeper readjustments of the system's classifications after prolonged exposure of the system to more sustained experience of anomaly.

- Underdetermination of theory by observation is a consequence of the fact that the assimilation of empirical data is only part of the story in the systemic knowledge-generating process. Even given that empirical adequacy is a major contributor to the perceived usefulness of scientific contributions and therefore an important factor in the assimilation process, the assimilation has to proceed, not simply by scientists choosing to incorporate new empirical results but through the progressive rearrangements of transaction patterns that such choices initiate. Since these rearrangements are rearrangements of an existing structure, it is most likely that piecemeal, local adjustments would take place. The resulting global classifications of the system would show only minor change, giving the appearance of a Duhemian reaction to new data.

- Incommensurability of past and present scientific knowledge is not surprising in the face of the major rearrangements of the classificatory categories which can result, after a prolonged assimilation process, from the incorporation of significant new inputs. The extent of the incommensurability will depend on the extent of the differences between the past and present classifications, a difference which is likely to grow as more time passes and the system undergoes further adaptation. The case of the incommensurability of currently competing theories is also a matter of the difficulty of making comparisons from points of view based on different classifications, only here the classifications are more local – representing different schools of thought within science.

- Theory-ladenness of observation is the science analog of what Butos and Koppl (2007: 29-30) have called the *pons asinorum* for *The Sensory Order* (i.e., Hayek's insight that memory is antecedent to sensation) applied to science rather than to the brain. An adaptive system's reactions to input depend on *both* the input and the operative classifications, where the latter are the contingent result of the system's past experience and determine not only how current observations are to be interpreted but also which observations are to be considered relevant.

In short, viewed as systemic attributes of knowledge, underdetermination, incommensurability, and theory-ladenness are all natural consequences of the way in which the differentially mutable structure of an adaptive system reacts to its environment.

We can summarize the epistemological implications of adaptive systems theory as follows. Scientific knowledge is not just a selection from the best offerings of individual scientists; it is the outcome of an extended, institutionalized, social process of publication, criticism, interpretation, citation, argumentation, promotion, rejection, reinterpretation, assimilation, and even (from some participants' points of view) misinterpretation of those individual contributions. It is not inherently propositional (although aspects of it can be usefully expressed in such terms) and is not any simple summation of individual contributions chosen for inclusion by participating scientists (although individual contributions are obviously essential elements in the system's classification process). Knowledge generation in science cannot be separated from the procedures through which scientists interact and through which individual efforts are transmuted to become accepted as scientific knowledge, and to attempt to characterize "scientific knowledge" apart from "scientific process and institutions" is a dualism on par with separating mind from matter. It is not simply that the norms of science, having emerged out of past interactions between scientists, have a determining effect on which individual contributions are acceptable contributions to knowledge – it is that such norms, procedural as well as prescriptive, are themselves part of the process through which individual scientific contributions are transmuted into emergent scientific knowledge. Not all social effects are simply influences on individuals on a par with personal interests; the particular transactional forms involved in the absorption of individual contributions are an integral part of the process of knowledge generation. Individuals and their scientific interactions together form a networked system, and scientific knowledge is an emergent side-effect of the activity within that system. To be more specific, scientific knowledge is a classification by the system of science of the environment to which that system is sensitive and to which it adapts – indeed, "knowledge" in general should refer to the classification produced as a side effect of the adaptation to its environment by *any* adaptive system.

Knowledge, thus construed, does not exist apart from particular knowledge-generating systems, and the capacity for knowledge is determined by the structure of the system. The individual brain performs a classification on events of the external world detected by the sense organs – a classification that Hayek (1952) characterizes in terms of "rules of action" that are the propensities for responses to these events. Individual knowledge *is* this classification. Hayek (1978: 40) is clear that "it is meaningless to speak of perceiving or thinking except as a function of an acting organism in which the differentiation of stimuli manifests itself in the differences of the dispositions to act which they evoke". Analogously, the sort of knowledge generated by a market order, or a scientific order, or even a political order, being the result of external events impinging on those complex objects undergoing a process of categorization unique to the particular type of order involved, depends both on the sorts of events being classified and the internal structure of the system doing the classification.

We are so used to thinking of knowledge as something "out there" to be tapped into as best we can that its necessary connection to the system in which it is generated, and the fact that it cannot be talked about sensibly independently of that system, is easily overlooked. Even Hayek, especially in his earlier work, seems to overlook it, as, for example, when (1945: 520) he speaks of "the utilization of knowledge not given to anyone in its totality". This way of putting the matter fosters the notion that there actually is such a thing as that totality, and that the problem is how best to access it. This can be contrasted with his later statement quoted above (1978: 40) emphasizing the contingency of individual knowledge with respect to the structure in which it is generated, and, perhaps even more pointedly, with his statement (1952: 4) that "a question like 'what is X?' has meaning only within a given order". The lack of clarity evident from the juxtaposition of these appraisals of the status of knowledge has its root in the failure in the 1945 statement to distinguish between different sorts of order-dependent knowledge. Individual people know "the particular circumstances of time and place" (knowledge generated in the context of a sensory order), but this is not knowledge in the context of a market order. Individuals deploying their individual knowledge with respect to useful goods and services in the interactive context of a market generate knowledge in the market sense – market prices, reputations, and brand names. Individual knowledge is not accumulated and made manifest at a market level in this process; it obviously conditions the individual interactions, but what emerges as a side-effect of these interactions is something entirely different – not a summary of individual knowledge, but a transformation of it which can be seen as a classification of events in the market system's environment. From the point of view of individuals, this market knowledge feeds back as inputs from which enhanced individual knowledge can be generated. The market, however, does not make evident those particular circumstances of time and place any more than the sensory order makes evident the events in the external world as they "really are" – it classifies them, produces an abstract representation of them, in a form that individuals can use to their advantage by generating new propensities for action and thereby adapting more successfully to their own environment.

Learning from this confusion over the identification of knowledge, we can say that the pursuit of a study of the knowledge-generating class of objects called adaptive systems should begin by being clear, for any particular type of system, exactly what is being classified, what categories are being produced by this classification, and how is it that this classification is taking place. The "how" involves elucidating the structure behind the order – of what elements is it built, what are the relevant characteristics of these elements for the functioning of the structure, how do the elements interact, and how does this interaction give rise to a classificatory apparatus. Such an approach would have epistemology become more like a positive natural science and less of a branch of normative philosophy.

8. Conclusion

Generalizing from Hayek's description of the fundamental structure and operation of the neural order, we have shown that certain social orders (markets and science, specifically) can be regarded as operating according to similar abstract principles, although the implementations of these principles differ considerably in the two cases. The common picture that emerges is as follows:

1. The system must be organized as a definite but mutable structure, the analog of Hayek's "map". In the neural order, this structure is completely physical, but what characterizes social orders is the capacity for the structure to have, in addition, more abstract elements – institutions, for example. This structure can be changed by the effects of the interactions going on within it, and some parts of the structure can change more rapidly than others. In the market, the structure is comprised of the institutional framework, the market participants and their routines, the particular goods transacted, and their market prices. In science, the structure is comprised of the institutional framework, the arrangement of fields and their publication outlets, the accepted base of scientific knowledge, and the reputations of scientists. The structure in all cases contains interactive components between which transmissions can occur (neurons in the neural order and people in social orders) that affect the components and may change their subsequent activity.
2. There is an ongoing pattern of transmissions between the components of the system, the analog of Hayek's "model". This pattern is sensitive to incoming stimuli and can induce effects on the environment. Environmental changes trigger transmissions between system components (or changes to ongoing transmissions) that (a) produce follow-on transmissions within the system which follow the (possibly re-entrant) paths inherent in the existing structure, and (b) can result in alteration of elements of the map structure itself – especially the more mutable parts but, in principle, anywhere.
3. This dynamic arrangement functions as a classifier of stimuli, and through this process of classification it is able to learn to adapt to its environment – to effectively generate knowledge of its environment. In social orders, the concrete effects of the classification process on the map (market prices or scientific reputations, for example) are observable by the participants and can induce follow-on responses which have effects analogous to the inducement of patterns characteristic of "expected" stimuli that occurs in the neural order.

The claim that social arrangements (markets and science of course, but quite possibly others, including city neighborhoods and market firms) can be regarded as systems capable of adapting to their environments is, by itself, vague and insubstantial, and not productive of much insight. But, if it is augmented with a theoretical blueprint of how adaptive systems in general must be structured and how the consequences of such structure are empirically manifest, it becomes a much more incisive assumption, allowing one to identify and observe the basic principles of organization and operation in a range of different situations. By highlighting generic similarities across a range of social arrangements, this approach has the capacity to bring together a great deal of existing work that has proceeded independently, and to point the way to further study stemming from the recognition of this commonality of structure.

Acknowledgements: We are grateful to the participants at the Second Conference on Emergent Order and Society (Portsmouth, NH, November 1-4, 2008) for helpful discussion, and we thank Gus diZerega for his critical comments. Thomas McQuade thanks Mario Rizzo for providing institutional support.

Notes

[1] Hayek's somewhat idiosyncratic employment of the terms "map" and "model" (which we adhere to throughout this paper) may mislead the reader into thinking of them as entities subject to conscious perusal. They are not. They are descriptive of the structure and functioning of the network, one emergent effect of which (in the case of brains, and some brains only, as far as we know) is consciousness.

[2] We do not wish to underplay the significant differences between brains and social arrangements. A major one is that neurons do not act telelogically but the system in which they function does; people have purposes but the system in which they function doesn't. For neurons, chemical and electrical factors are involved in interactions and their rearrangements; for people, teleological factors play an important role in interactions and their rearrangements. But, being concerned simply with the adaptive properties of these structures viewed as generalized re-entrant networks, we can abstract from the *specific* ways in which interactions and their rearrangements are implemented.

[3] This is not intended to be an exhaustive taxonomy of groups. For example, bureaucratic organizations would be an obvious third category. We emphasize this particular distinction simply because groups which function as adaptive systems are often lumped together with those which function additively.

[4] These Kirznerian terms ("stifled" and "superfluous") used by Ikeda have a normative ring to them, but whatever value-judgments are implied or imputed are no part of adaptive systems theory. The phenomena being described are real enough, and these are what are of interest.

[5] About two-thirds of the funding of academic scientists in the U.S. comes from federal and state government. While such funding supports only about 10% of all R&D in the U.S., it supports in excess of 50% of basic research, according to the NSB (2006: 5-12).

References

Barnes, Barry. 1974. *Scientific Knowledge and Sociological Theory*. London: Routledge & Keegan Paul.

Benson, Bruce L. 1989. "The Spontaneous Evolution of Commercial Law," *Southern Economic Journal,* 55(3): 644-661.

Bloor, David. 1976. *Knowledge and Social Imagery*. New York: Routledge & Kegan Paul.

Boehm, Stephan. 1994. "Spontaneous Order," in Geoffrey M. Hodgson, Warren N. Samuels, and Marc R. Tool eds, *The Elgar Companion to Evolutionary and Institutional Economics: L-Z*, (Aldershort: Edward Elgar).

Buckley, Walter. 1998. *Society – A Complex Adaptive System*. Amsterdam: Gordon and Breach.

Butos, William N. 2003. "Knowledge Questions: Hayek, Keynes and Beyond," *The Review of Austrian Economics*, 16(4): 291-307.

Butos, William N. and Roger Koppl. 2007. "Does *The Sensory Order* Have a Useful Economic Future?," *Advances in Austrian Economics* 9: 19-50.

Butos, William N. and Thomas J. McQuade. 2002. "Mind, Market, and Institutions: The Knowledge Problem in Hayek's Thought," in J. Birner, P. Garrouste, and T. Aimar eds, *F.A. Hayek as a Political Economist*, (London: Routledge) 113.

Butos, William N. and Thomas J. McQuade. 2006. "Government and Science: A Dangerous Liaison?," *The Independent Review* 11(2): 177-208.

Butos, William N. and Thomas J. McQuade. 2009. "Nonneutralities in Science Funding: Direction, Destabilization, and Distortion," unpublished manuscript.

Cheung, Steven N.S. 1975. "Roofs or Stars: Stated Intents and Actual Effects of a Rents Ordinance," *Economic Inquiry* 13: 1-21.

Churchland, Paul. 1995. *The Engine of Reason, the Seat of the Soul: A Philosophical Journey into the Brain*. Cambridge, MA: MIT Press.

Coats, R. Morris and Gary M. Pecquet. 2008. "The effect of opening up ANWR to drilling on the current price of oil," unpublished manuscript., accessible at *Munich Personal*

RePEc Archive, http://mpra.ub.uni-muenchen.de/9543.

Goldman, Alvin. 1987. "Foundations of Social Epistemics," *Synthese* 73: 109-144.

Goldman, Alvin. 1999. *Knowledge in a Social World.* Oxford: Oxford University Press.

Hayek, F. A. 1945. "The Use of Knowledge in Society," *American Economic Review* 35(4): 519-530.

Hayek, F. A. 1952. *The Sensory Order.* Chicago: University of Chicago Press.

Hayek, F. A. 1973. *Law, Legislation and Liberty: Vol. 1, Rules and Order,* Chicago: University of Chicago Press.

Hayek, Friedrich A. 1978. *New Studies*. New York: Routledge & Kegan Paul.

Hull, David L. 1988. *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science*. Chicago: University of Chicago Press.

Ikeda, Sanford. 1997. *Dynamics of the Mixed Economy*. New York: Routledge.

Kauffman, Stuart A. 1993. *The Origins of Order: Self-Organization and Selection in Evolution*. New York: Oxford University Press.

Kevles, Daniel J. 1998. *The Baltimore Case*. New York: W.W. Norton.

Kitcher, Philip. 1993. *The Advancement of Science*. Oxford: Oxford University Press.

Koppl, Roger. 2002. *Big Players and the Economic Theory of Expectations*. New York: Palgrave Macmillan.

Kuhn, Thomas S. 1962. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

Laudan, Larry. 1990. "Normative Naturalism," *Philosophy of Science* 57: 44-59.

McQuade, Thomas J. 2000. *Orders and Institutions as Objects of Investigation*, unpublished Ph.D. dissertation, Auburn University.

McQuade, Thomas J. 2007. "Science and Market as Adaptive Classifying Systems," *Advances in Austrian Economics* 9: 51-86.

McQuade, Thomas J. 2009. "Science and The Sensory Order," *Advances in Austrian Economics* 11 (forthcoming).

McQuade, Thomas J. and William N. Butos. 2003. "Order-Dependent Knowledge and the Economics of Science," *The Review of Austrian Economics,* 16(2): 133-152.

McQuade, Thomas J. and William N. Butos. 2005. "The Sensory Order and Other Adaptive Classifying Systems," *Journal of Bioeconomics* 7(3): 335-358.

Merton, Robert K. 1973. *The Sociology of Science.* Chicago: University of Chicago Press.

Mises, Ludwig. 1949. *Human Action.* Chicago: Henry Regnery.

National Science Board. 2006. *Science and Engineering Indicators 2006, v.2, Appendix Tables.* Washington: National Science Foundation.

Padgett, John F. 2001. "Organizational Genesis, Identity, and Control: The Transformation of Banking in Renaissance Florence" in James E. Rauch and Alessandra Casella eds, *Markets and Networks*, (New York: Russell Sage) 211.

Padgett, John F. and Paul D. McLean. 2002. "Economic and Social Exchange in Renaissance Florence," *Santa Fe Institute Working Paper* 02-07-032.

Pickering, Andrew. 1984. *Constructing Quarks*. Chicago: University of Chicago Press.

Quine, William V. 1969. "Epistemology Naturalized," in *Ontological Relativity and Other Essays*, (New York: Columbia University Press) 69.