Journal of Bioeconomics (2005) 7:335–358
DOI 10.1007/s10818-005-7606-5 © Springer 2006

The Sensory Order and Other Adaptive Classifying Systems

THOMAS J. MCQUADE
Department of Economics, New York University, New York, NY 10012, USA (thomas@mcquades.net)

WILLIAM N. BUTOS
Department of Economics, Trinity College, Hartford, CT 06106, USA (william.butos@trincoll.edu)

Synopsis: Hayek's cognitive theory, which seeks to describe the operation of a particular order, in fact provides a paradigmatic account of knowledge-generating orders in general. We claim that this paradigm provides a fertile conceptual framework for exploring a variety of problems in economics and social theory. In particular, we shall show that Hayek's conception of the 'map' and the 'model', which he uses to explain the operation of the complex adaptive classifying system called 'mind', are promising analytical devices with applications extending to social structures of various kinds and complexity. We use Hayek's notion of the map and model to analyze how different social structures – regarded as classifying systems – work in terms of their input, processing, and output capabilities. The adaptive characteristics of such systems, via communicative routines, multi-level classification, and feedback, form central motifs for our discussion of markets, science, and other social structures. We show that by analyzing the knowledge-generating characteristics of such structures we are also able to gain insights about the circumstances affecting their adaptive properties.

Key words: social theory, cognition, knowledge, adaptation, markets, science, classification, Hayek

JEL classification: A12, B41, B52, B53, Z13

1. Introduction

Hayek's (1952) *The Sensory Order* describes an 'adaptive classifying system' – a system which adapts to its environment by forming an internally maintained model of that environment. In very general terms, he characterizes the structure of that system as implementing a mutable arrangement which he calls the 'map' because of its ability to represent the general features of its environment. He shows how such a structure has the potential for maintaining and continuously updating a model of the environment as currently experienced – a model that is capable of anticipating expected changes in that environment. The map, as Hayek explains, can be seen as an apparatus capable of performing a multi-level classification of the input stimuli experienced through the sensory apparatus of the system, including stimuli, both external and internal, resulting from actions of the system on the environment itself. The map and its model constitute the system's knowledge of its environment, enabling stimuli to be classified against past experience and responses to be selected that take into account the success or otherwise of past responses to similar stimuli. It is not so much a 'store' of knowledge as an ongoing, self-modifying and self-organizing process, whose interactions with its environment continually update its classificatory ability, refining existing classes, and introducing new ones as experience is accumulated.

While Hayek's work has had little direct influence on modern neuroscience, it is not incompatible with the more mainstream developments from Hebb (1949) to Edelman (1987, 1992). And it was pitched at a level of abstraction that, while assuming some basic properties of neurons and their connections, did not require the level of anatomical sophistication available today. But its distance from biological detail, while perhaps seeming like a drawback to a neuroscientist, is something of an advantage for a social scientist – for Hayek, both directly and by implication, has focused on significant attributes of the system and its components, attributes that are the *sine qua non* for *any* system to function as an adaptive classifier of the stimuli received from its environment.

We propose to highlight these attributes, and then to show that their analogs are present, in different physical realizations, in various social systems. This identification enables us to characterize these social arrangements – which, specifically, include markets and science – as 'adaptive classifying systems' which build up adaptive models of their environments based on stimuli and feedback from those environments and which therefore can be said to generate, each in its own unique way, 'knowledge' of those environments. The adaptive capabilities of such systems, which involve communicative routines, multi-level classification, and feedback, form central motifs for our discussion of markets and science. By focusing on the requirements for such knowledge-generating processes to function successfully, we are able to point to circumstances that are likely to impede or short-circuit such functioning. We suggest that the analysis of the knowledge-generating characteristics of these social structures, and others such as city neighborhoods and firms, constitutes an interesting and novel approach to the understanding of the functioning of such structures and of their adaptive properties.

2. Generalizable attributes of the Hayekian neural order

In *The Sensory Order*, Hayek provided a high-level description of the brain as a physical system capable of maintaining an order of sensory qualities by operating as an apparatus of classification. Although Hayek explicitly considered this one of his most important works,[1] it initially received very little attention from Hayekian economists and social theorists. But, with the rise in interest in cognitive issues in social theory over the past ten years or so, the work has started to find an audience, and capsule presentations of its content can be found in several papers.[2] Our emphasis here is, however, somewhat different from that of the authors cited in that the goal is generalization rather than exposition. To this end, the following digest of Hayek's (1952) description (augmented by a little biological detail not stressed by him), focuses on what seem to be the basic general characteristics of this particular classifying system, its components, and its mode of operation, and presents them with an eye toward showing that their analogs are also present in (and important to) various orders in the social domain.

The basic structure of the system is in the form of a network of components, physically connected, that interact via the transmission of electrical impulses. The components are neurons, cells capable of generating outgoing impulses ('firing') if stimulated sufficiently by incoming impulses, but limited in their ability to continuously fire by the requirement for a relaxation period between firings and because of the onset of fatigue. Each neuron is connected to many others (but nowhere near to all of the others in the network) by means of extensions called axons, and new connections can be established or existing ones eliminated depending on ongoing activity so that pathways with impulses often occurring concurrently will tend to become connected. The strengths of established connections (their ability to induce follow-on firing) can change depending on the activity they experience. The connections can form complex patterns in the network that involve groupings of neurons in which all neurons in a group connect mostly with the same downstream groups – but not exclusively, allowing for the possibility of re-entrant patterns that support the operation of feedback loops.

The system is open to its environment and interacts with it in both sensory and motor modes. Events in the environment (including those in the body of the organism containing the brain) impact on sensory organs that transmit impulses to the brain, and some neuronal pathways within the brain connect to motor neurons that can induce actions affecting the environment. The firings and subsequent refirings of interconnected neurons within the brain, induced by the flux of incoming stimuli, result in the system as a whole maintaining a continuous, fluctuating level of activity which forms a background over which reverberations due to newly incident stimuli are superimposed.

The network of connections (of varying strengths) between the component neurons, which changes as a long-term result of the stimulus-induced patterns of activity, effectively functions as a 'map' of the environment as experienced in the past in that it enables the emergence, from a given stimulus, of an induced pattern of impulses characteristic of that stimulus and of other potential stimuli which have accompanied the given stimulus in the past. This pattern of impulses generated in the map by the current stimuli is, therefore, a 'model' of the environment as currently experienced. The model is anticipatory in that trains of impulses originating from a stimulus induce a 'following', i.e., impulses in neuronal pathways that, in the past, have tended to be active at the same time as the pathways excited by the current stimulus (and these followings can in turn induce further followings). The model thus embodies expectations of subsequent stimuli by virtue of the fact that the map enables the inducement of patterns characteristic of associated stimuli even when those stimuli may not be actually present. Since connections exist to motor neurons at many levels and these connections, like all connections in the map, have been developed as a result of experience (phylogenetic or ontogenetic), the model can result in the selection of motor activity consistent with expectations it embodies as to the character of the current situation.

The network of neurons and their axons acts as a system of classification in that the impulses from stimuli produce particular patterns of follow-on impulses within the network. Two stimuli are classified as identical if they produce the same impulse pattern; they are similar to the extent that their patterns overlap. As both subject and object of classification are patterns of impulses, further classifications can proceed at higher levels in terms of the effects the (classified) follow-on patterns have on subsequent neuronal groups. This classification is multiple in several senses – any particular stimulus can be a member of multiple classes, an assignment of a particular stimulus to a class may change depending on the presence of concurrent stimuli, and classes (being represented in terms of impulses) can themselves be further classified at subsequent levels.

The map is effectively an apparatus of long-term memory in that its configuration enables the effects of one stimulus to reconstruct the effects of other stimuli with which it has been associated in past experience. (There is also short-term memory in that the effects of stimuli reverberate in the network for some time before being damped out.) The physical changes in neurons and their connections implement a form of learning in which the system adapts to its environment.[3] The network adjusts to injury to or loss of a component or connection (or to a localized group of components and connections) simply by continuing operation. The system as a whole will continue to function, but it will have 'forgotten' whatever long-term memory was embodied in the lost or degraded connections and its classification of stimuli whose followings are seriously affected by the loss will be changed (probably in unpredictable ways). The learning activity of the system will continue, however, and eventually the injury will tend to be 'worked around,' with other connections forming or being co-opted to perform classifications similar to those in effect prior to the injury.

This description characterizes the brain[4] as a system interacting with and adapting to its environment by performing a multi-level classification on the stimuli it receives from the environment – an adaptive system with no central locus of control. It has been pointed out by others[5] that there exist a variety of 'complex adaptive systems' besides the brain that seem to have some level of capability for adaptation through changes in structure based on system experience, but we propose to follow up that observation in a novel way by abstracting from the particular implementation of the neural order described by Hayek what seem to be the essential attributes of such a system. The following appear to us to be the necessary elements:

(1) *The structure of the system must implement an adaptive map*, changing as a result of activity induced within it by environmental stimuli and by stimuli originating from within the system itself in ways that, given the mix of

excitatory and inhibitory impulses its components generate, tend to maintain its general integrity while conforming it to the external environment and possibly acting on that environment. The system must contain reactive components, each component connected in some way to many others with the possibility of re-entrant loops. There may be identifiable groupings of components for which members have in common, more or less, the other groups to which they connect. The characteristics of the map itself – including the firing properties of the reactive components, the connections they make to other components, and the strengths of these connections – can change as a result of the activity within the system. The structure of the system is, therefore, mutable, and is conditioned by its experience.

(2) *The activity within the system must implement an expectational model*, driven by the effects of current stimuli on the map and capable of inducing responses on the environment consistent with the current map, responses which may, in some cases, anticipate events. Connected components will affect each other by transmission of an impulse. The propensity of a component to transmit an impulse depends on its own reception of impulses and on its own internal state. An impulse will not be transmitted unless the incoming impulses induce the component to exceed its current threshold for transmission. Some incoming impulses may be inhibitory in that they tend to reduce, rather than increase, the likelihood that the component will fire. Components cannot transmit continuously; there are limitations that considerably constrain the ability to transmit and thereby protect the system from chaotic overloading. While the effects of the transmissions from any group of components may propagate throughout the system, there is no group that is identifiable as a central locus of control in that it alone determines the transmissions of all other groups. The follow-on effects of stimuli (although tending to damp out as not all impulses induce a follow-on firing) will be projected on a background of other recent and concurrent stimuli, so that the overall pattern induced by a particular stimulus is characteristic not only of itself, but of what else is sensed by the system at the same time. This model of the environment may be expectational in that patterns resulting from current stimuli may evoke follow-on reactions within the system that dispose it to react as if to anticipate subsequent stimuli that have been associated with the original stimulus in the past.

(3) *The reactive components of the system must be sufficiently interconnected to support the ability of the system as a whole to classify stimuli.* The arrangement of map and model classifies in the sense that the patterns of transmissions induced by stimuli in the context of the reverberations of concurrent stimuli represent the attributes of the event inducing the stimulus that can be sensed by the system. Patterns of transmissions are at the same time both what is being classified and the results of the classification process. To the extent that the system can be decomposed into layers of component groups that act successively on transmission patterns, the classification is a multi-level one, and outgoing responses into the environment are possible at all levels. The classifications can be

> thought of, in general terms, as instantiating the system's knowledge of its environment. Injury to localized parts of the system may cause loss of the classifications supported by transmissions through the affected parts, but the classificatory process will continue, and the relearning process will tend to alter the map in ways that work around the injury.

Since we are dealing with structured but mutable systems[6] that adapt to their environments by maintaining and refining a classification of the stimuli they receive, we have chosen to label them as 'adaptive classifying systems'.[7]

3. On generalizing from neural to social orders

In focusing on some of the basic processes occurring in brains and proposing that they appear (in different implementations) in various social systems, we are both inspired and warned by Hayek's own discussion[8] of the differences between brains and social systems. First, the inspiration: 'There is, therefore, no reason why a polycentric order [like a social order] in which each element is guided only by rules and receives no orders from a centre should not be capable of bringing about as complex and apparently as "purposive" an adaptation to circumstances as could be produced in a system [like an organism with a brain] where a part is set aside to preform such an order on an analogue or model before it is put into execution by the larger structure' (1967, p. 74).

Then the warning: 'Although the brain may be organized on principles similar to those on which a society is organized, society is not a brain and must not be represented as a sort of super-brain, because in it the acting parts and those between which the relations determining the structure are established are the same, and the ordering task is not deputized to any part in which a model is preformed' (p. 74). We obviously don't think that Hayek's warning closes the door on the generalizing move we want to make. But it is worthwhile to analyze the warning closely, for it expresses, indirectly but concisely, the immediate objection most economists and other social scientists would deploy.

Hayek points out, quite correctly, that the sensing and motor parts of a social order are the same elements – people – as those whose relations determine the adaptive structure. This is certainly a factor to be taken into account, but it is difficult to see why it should imply that the maintenance of a map and the support of a model cannot occur in social orders in ways that are similar, in principle, to the processes generating the sensory order. Further, Hayek's allusion to a sequential process – a model is preformed in the brain, then the appropriate motor action takes place – oversimplifies his own discussion (1952, pp. 84–85) in which classification and motor response proceed simultaneously in a multi-level, parallel process. This more complex picture of interleaved classification and response is quite compatible with an arrangement in which the components involved in classificatory interactions and the components involved in implementing responses may be physically the same components. Perhaps what Hayek is actually warning against is not the idea that societal orders can, at some level of abstraction, function in ways analogous to *brains*, but that they can usefully be modeled as functioning like *organisms with a brain*, or, worse, like organisms with a 'super-brain' cabal of philosopher-kings making all of the decisions and issuing all of the orders for the rest of society to follow mindlessly.

Let us, then, be very clear. We *do not* think that society in general, or any of the various social orders that make up society, is a brain, let alone a 'super-brain'. 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.[9] 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 that constitute adaptive reactions of the system to changes in its environment.

4. The market economy as an adaptive classifying system

The background for our characterization of the market economy as an adaptive classifying system is the conception of the market economy as catallaxy which animates the work not only of Austrian economists such as Mises and Hayek but of many others – Buchanan, for example – as well. A catallaxy is understood to comprise an order of interacting agents who engage in the voluntary exchange of claims to property (Mises 1966; Hayek 1975, 1976). Organized markets are an institutional arrangement that allows buyers and sellers, cross-sectionally and through time, to regularly exchange with each other (Buchanan 1979). Such catallactic activity generates a unique kind of output in the form of market prices. The separate activities of many individuals, each motivated by their respective goals and intentions, interact in ways that produce at every moment a constellation of market prices for goods and services – an outcome that could not have been generated in any other way (Boehm, 1994), since the catallaxy's capacity to produce such prices is an emergent property of the system and such outcomes, as

Buchanan & Vanberg (1991) remind us, cannot be defined independent of the process itself.

This conception of a market economy abjures the usage of a formal device such as the Walrasian auctioneer or an idealized central planner as a *deus ex machina* who ensures coordination, and points instead to the coordinating mechanisms, supported by evolved institutions of property, contract, and exchange, endogenous to the system itself. The pre-eminent characteristic of catallactic functioning is that the personal, often private, knowledge of each individual in the market undergoes a transformation as a consequence of catallactic interactions, producing as an emergent byproduct of those interactions marketable goods and market prices that tend to become consistent with the abilities and expressed individual preferences of market participants.

Such a picture of the market economy readily lends itself to description in terms of map and model. The market's analog of the map is a stable but mutable structure built from the following elements (listed in rough order of increasing mutability):

(1) The institutional framework circumscribing and stabilizing the nexus of market interactions. Market exchanges are subject to prevailing rules, in the sense of a 'statement by which a regularity of the conduct of individuals can be described' (Hayek 1967, p. 67), and to such constraints governing the use and disposal of property, including mechanisms to ensure the discharge of contractual agreements and to provide for remediation. These are the deep and typically long existing institutional structures without which it would be difficult to imagine catallactic activity occurring on any large scale.
(2) Those semi-fixed routines that transactors have learned to rely on to implement their plans. In the daily business of life, to slightly modify Marshall's felicitous phrase, people ordinarily undertake activities that follow certain patterns; they function within a corridor of activity that on a daily or weekly basis turns out to be broadly repetitive. While there may be a great deal of agitation in the particulars of activity, from a loftier perspective we see a kind of regularity in the contours of agent behavior. Because these routines, habits, and conventions are not ordinarily subject to sudden and dramatic change, they form important stabilizing elements in the nexus of market activity.
(3) The market participants themselves, whose behaviors form overlapping spheres of interactions directly and indirectly as buyers and sellers, and whose tastes and preferences can be changed as a direct result of market experience – an effect that tends to be ignored in textbook economic theory but which can feature prominently in the phenomenon of market clearing.[10]
(4) The structured constellation of marketable goods and services, related to each other both as inputs and outputs and as complements.
(5) The market prices of these goods.

The analog of the model is the ongoing flow of transactions (characterized by voluntary transfers of goods and observable exchange prices) between the market participants – the reactive components of the system. These transactions are induced by stimuli from outside the system, such as environmental conditions and resource endowments, and by stimuli from within the system itself, emanating from the preferences and reactions and creativity of the market participants themselves. In this context, it is important to be clear about what it means to say that the market receives stimuli. Unlike the neural order, the market as such does not have dedicated receptor organs. But it is comprised of numerous agents and it is through their *actions* in the marketplace that the analogs of the sensory stimuli of the neural order enter the system as inputs.[11] Whether stimuli operating on the market originate from outside or from within the system, all such stimuli pertain to definite actions agents are performing or involve a communicated intention to act in a particular way.[12] In any case, the transactions of the market order, and the effects those transactions have on the people who experience them, can readily change, in ways well known to economists, the spectrum of goods and their market prices and, probably over a somewhat longer time frame, the habitual transaction patterns, routines, and preferences of market participants.[13]

To illustrate how the market economy can be regarded as functioning in a classificatory and anticipatory manner, consider an exogenous environmental change that brings drought to the wheat-growing regions of the country, resulting in an anticipated reduction in the supply of wheat. This constitutes a change in the stimuli that will affect the price system, and these effects, together with the behavioral responses they trigger, are part and parcel of the economist's stock and trade and need not be dwelled upon here. However, in pointing to the intertemporal adjustments which the wheat situation is very likely to generate, it is possible to highlight the 'anticipatory' aspects of the market order. In the wheat market, commodities speculators bid up prices for both current inventories of wheat and future (drought-affected) anticipated supplies. This activity has the unintended side effect of generating a structure of prices in the wheat market (and in other related markets) that take account of the anticipated supply conditions through time. Without any centrally -administered directives, the market has channeled individual activity, including activity far removed from the local circumstances of wheat production, in ways that establish the conditions for consumption of wheat products to be 'smoothed out' during the drought.

The market has, in effect, functioned as an anticipatory mechanism triggered by a current stimulus (a drought in wheat-growing regions) that has conditioned it to respond in ways that take account of the future. It is 'as if' the system knew that the drought would exact its toll through the next harvest and had responded in a manner that was both adaptive to the change in circumstances and anticipatory of the future. Without being conscious and without any central locus of control, the system displays an emergent kind of behavior and generates particular configurations in the generation of knowledge, principally in the form of prices and quantities through time and space, that could not have been produced in any other way. While it is clear that part of the process involves regular human insight into the future, it is at the market level – in the changes in market prices and quantities traded in both spot and futures markets – that we observe a manifestation of the anticipatory aspects of the system as it responds to environmental change.[14]

In short, the stimulus-response properties of the market are describable as involving a semi-fixed map that channels and registers incoming stimuli and can be changed by the effects of that process. Given the flow of impulses, the system develops a model of its environment, and if the correspondence between the existing map and the stimulus-generated current model exerts no pressure for change on the prevailing map, then we can think of such situations as 'business as usual'. But if we throw into the mix an environmental change that reveals discordance between the current (and now altered) model and the pre-existing map, an adaptive response may be called for. The resulting changes in the map are likely to be local, marginal, continuous, and not perfectly predictable. Thus, it may not be possible to specify with a high degree of precision that point when the system (or some part of it) responds in a 'business *not* as usual' fashion or even the particular form such a response will take. All the same, the ongoing and nuanced interaction between stimuli, model, and map suggests adjustments occurring in the system driven by the effects of an altered model on the map. In effect, the system reaches a certain (though *a priori* unspecifiable) point at which the process of reclassification has superimposed additional elements on the map or otherwise revised the map, reflecting the system's new adaptive posture and preparing it to execute different responses in light of the changed environment.

For the market economy, these adjustment responses are instantiated by market mechanisms in which entrepreneurial activity and the role of time markets play decisive roles. While economists often (and confusingly) speak of such adjustment in an 'atomistic sense' as if the market were a single agent or governed by a single mind, the relevant point here is the fact that it can occur only within the particular institutionalized setting of *market interactions* in which purposeful agents – entrepreneurs – generate and act generally upon local circumstances. Thus, such actions in actuality constitute the flow of stimuli the market absorbs; and it is from these stimuli that changes ensue that generate market-level phenomena, most obviously in the form of changes in the market prices in both spot and, as we noted earlier, in futures markets. This generation of market-level phenomena is not a process by which stimuli are simply aggregated or reconfigured, but a *transformation* into something different. The actions of the individuals affected by an exogenous event (e.g., the previously discussed drought) comprise the analog of sensory stimuli to the market order which, from our vantage point as market observers, produces a model of the current environment different from that existing prior to the event. Given the overlapping spheres of interaction among market participants, these effects and the resulting stimuli they induce become increasingly strong and pervasive. We can characterize the situation as one where the disparity between the model and the existing map becomes too great to maintain: in effect, the behavior of agents and the current constellation of market prices and outputs do not square.

Economists may speak of this as a manifestation of 'plan discoordination' and point to the springing into action of profit-seeking entrepreneurs, which is fine up to a point. But such descriptions do not adequately capture the 'sensory order' perspective proposed here. The process we wish to highlight is that of the continuous and incremental market generation of new knowledge in the form of different market prices for products. In short, the map constructed by the market changes. What makes this process difficult to capture analytically is that this new knowledge also serves as relevant inputs for agents that will alter behaviors yet again. The process is relentlessly open-ended, involving both negative and positive feedback, trial and error, and experimentation that continues to generate new stimuli and new market-level responses until, in the case considered here, the full effects of the drought have been internalized by the market and its components or, in sensory order terms, the map and model of the market are sufficiently compatible.

As a general rule, these sorts of adjustment processes in the market, which we have characterized in terms of map and model, occur quite seamlessly as if, indeed, they were silently guided by Adam Smith's metaphorical 'invisible hand'. And, in the example above, the anticipatory response of the system gets played out as a series of *unintended and unanticipated outcomes from the perspective of the individual components*. As a more complex example of this most important feature of market functioning, consider the imposition of maximum price controls in the apartment rental market of a large city. While the intention is to keep rents low, the price control's effect is to constrain the market's ability to generate market prices consistent with the kinds of interactions that would normally prevail. It is as if certain pathways in the market's map have become paralyzed or disabled. The current model in the damaged map exerts pressure on the map to change. The specific responses often observed include degradation in the quality of rental units, lower turnover among tenants in rent-controlled apartments, and the emergence of a variety of new institutions, such as 'key money' and 'furnished' apartment fees. The long-term effects may include a depleted stock of low-income rental units and a disproportionate increase in middle- and upper-income housing as entrepreneurs respond to perceived opportunities in sub-sectors of the housing market. Such responses represent induced workarounds in the market's map that formed no part of the intentions of those who initiated the price control.

In summary, we characterize the market order's functioning as resulting in the generation of a classification – most obviously, a spectrum of goods and their market prices – over a set of inputs, and therefore we claim that the principal property of the order is its knowledge-generating capacity (Butos & McQuade, 2002). The market is a complexly organized structure with no central locus of control. It functions as an evolving and adaptive order with specialized mechanisms of positive and negative feedback, particularly the price system, which work to align the overall operation of the order with exogenous environmental changes and constraints and with endogenous changes produced by the system and its components. That the market economy appears to adapt with greater agility and resourcefulness than the available alternatives (such as central planning) is, we suggest, precisely attributable to its superior capacity as an adaptive classifying system.

Detailing the *anticipatory* responses of the market via interactions between map and model appears as an important payoff to the approach discussed here. The inherent dynamism of the map-model analogy and of the market as an adaptive classifying system point away from static economic theory and toward analyses for which the time element is decisive, such as business cycle theory, monetary economics, entrepreneurial theory, institutional change, and economic growth.

5. Science as an adaptive classifying system

The background for our characterization of the scientific community as an adaptive classifying system is the conception of 'science as a process' described by Hull (1988) and elaborated by McQuade & Butos (2003). By 'science' we mean the activities of individuals whose interactions involve specifiable modes of behavior or routines that produce, as a byproduct of those interactions, a particular kind of emergent phenomenon we call 'reliable and codifiable' knowledge about the real world. We view science as a complex order comprised of components (individual scientists) who define and pursue their own aims[15] and whose behavior is situated and constrained within an institutional structure of routines, conventions, and organized groupings that are themselves ordinarily and substantially emergent phenomena of the order itself. There is no presumption that the components of the order have consciously selected or chosen to follow scientific procedures – like much of ordinary life, the rules and conventions of science are often tacit, internalized, and followed blindly. Given this backdrop, the principal conduit of interaction among scientists involves 'communication acts' of various kinds, primarily publication, but also including lectures, speeches, conference papers, and other vehicles by which personal (and necessarily dispersed) knowledge or ideas and analytical skills and techniques are transmitted within the system. In the process, the personal knowledge of individuals is transformed into order-generated scientific knowledge, although, of course, there is no way to predict in advance which particular individual inputs will actually contribute to scientific knowledge. This suggests that the emergent characteristics of the order, like that of the catallaxy, are defined in the course of the process itself.

Again, such a picture of the scientific community readily lends itself to description in terms of map and model. Science's analog of the map is a stable but mutable structure built from the following elements (listed in rough order of increasing mutability):

(1) The institutional framework circumscribing and stabilizing the nexus of interactions between scientists. These include the norms of publication and citation which, although they appear to have taken their current form as recently as

the late 1600s,[16] are the critical institutional structures without which it would be difficult to imagine scientific activity occurring on any large scale.
(2) Those semi-fixed routines that individual scientists have learned to rely on to implement their plans. This includes their organization into schools and groups, and their patterns in selecting their usual outlets for communication.
(3) The scientists themselves, whose behaviors form overlapping spheres of interactions directly and indirectly as writers, referees, and users of published ideas, and whose skills, tastes, and preferences can be changed as a direct result of their experience as scientific researchers.
(4) The reputations of individual scientists.
(5) The body of established scientific knowledge.

The analog of the model is the ongoing flow of transactions (characterized by publication in various forums and citation when invoking the work of others) between scientists – the reactive components of the system. As with the market, the stimuli which induce the transactions can originate either inside or outside the system.

In applying the map–model analogy to science, we first recognize that, like cognitive activity, science generates a particular kind of classification that describes aspects of the real world. We see science as a social activity driven by a particular institutional arrangement that produces as a byproduct of its operation a classification over a range of inputs – the classification called 'scientific knowledge' – that is unique to the circumstances of its generation. We also note that science, like the neural order and the catallaxy, has no central or conscious locus of control in that there is no single element or group that directs the operation of the overall order, although in the social realm the vulnerability of such orders to constructivist influence makes them susceptible to just that. There are several potential kinds of costs associated with attempting to turn, in Hayek's (1973, ch. 2) terminology, a cosmos into a taxis (or a decentralized spontaneous order into an hierarchical organization). Certainly, for the discussion here, they would include a reduction in the order's receptivity to its environment and additional top-down constraints on the kinds of connections its components can establish among themselves as the imposition from above of centralization and conscious control increases.

When scientific activity functions as an emergent order, however, the panorama of diverse and novel stimuli arising from its components and the possibility that such stimuli have the capacity to intermingle with other components allow the overall pattern to exhibit the effects of more than just individual stimuli. In science these stimuli become subject to the deployment of various specialized feedback mechanisms that sift through the welter of inputs, determining which rival theories and explanations, i.e., classifications, will survive. Science, then, has the capacity to weed out the *relatively* weaker claims put forward by scientists and it does so by critically evaluating and rejecting (if only provisionally) multiple and rivalrous theories. Clearly, the greater the degree of imposed centralization, the more limited will be the flow of stimuli from the components of the order and the fewer the number of rival theories that will surface. In the limit, if the analog of economic central planning was applied to science, the resulting structure and its emergent characteristics would more closely resemble that of a single mind. Learning in such a system would be constrained to what that single mind could learn; the imposed institutional arrangements would effectively shut down the kind of learning that is a byproduct of interactions among individuals.

Within the nexus of the institutional arrangements of the scientific order, we can suggest at a suitably general level of the analysis that individual scientists produce stimuli that reflect the autonomous creativity of their own cognitive mechanisms in the context of what they have learned through study, tuition, and instruction, i.e., the 'background' canvas of order-generated knowledge that has survived various kinds of filtering mechanisms or critical scrutiny and, hence, qualifies as reliable and codifiable knowledge. We may think of this as 'received scientific knowledge', although we are not suggesting that scientific knowledge is monolithic. Indeed, it is often difficult in certain disciplines to empirically ascertain lines of demarcation concerning the status of knowledge, especially when, as often is the case, different and possibly incompatible theoretical visions or disputed empirical findings dominate. Such situations, in fact, actually highlight the complexity of the scientific order, suggesting that the various sub-orders found in science and across disciplines have evolved in ways that have produced highly specialized institutional arrangements. But notwithstanding the particular emergent characteristics different branches of science or different perspectives within branches may have generated, each branch or sub-branch, nonetheless, presents to its practitioners a body of knowledge which provides a kind of starting point for their own activities.

This body of received knowledge, including tacit knowledge such as laboratory skills and certain kinds of deeply internalized learned behavior, is a major component of the scientific order's analog of the map. The functional significance of the map in science is that it constitutes the order's 'long-term' memory, hence serving as the backdrop or context for newly generated ideas. The map, in effect, is the institutional manifestation of a 'not anything goes' constraint that helps to maintain continuity and stability within the order by retaining knowledge that has already survived some critical process and thereby proven its value to the order's adaptation to its environment. But the map in science, just as its analogs in the neural and market orders, is not static, but of necessity subject to ongoing changes, typically small but sometimes identifiable even over relatively short periods of time. It is the mutability of the map in science in response to the continuous flow of stimuli coursing through the various media available to scientists that enables the order to modify or generate new classifications and thereby establish a closer correspondence between its output and its environment (i.e., the generation of scientific knowledge and its domain of inquiry).

Can there be any question of the success of science in generating new knowledge? To account for this growth of knowledge, we think that certain interesting elements in what is surely a very complex story can be broached by invoking the map–model characterization. The institutional routines and communication channels of science help to ensure that the conjectures, theories, and empirical evidence generated by scientists are largely compatible with the existing pathways of the map, i.e., with accepted scientific knowledge. At the same time, however, portions of the incoming stimuli may form an interpretation of the environment that requires some kind of modification on the existing classification instantiated by the current map. Such stimuli in science may be unexpected empirical results or different theoretical hypotheses that dispose the order to be sensitive to further stimuli, both corroborative and otherwise. This kind of anticipatory role of the model highlights the capacity of the system to essentially sensitize itself to new stimuli, perhaps preemptively discounting certain kinds of stimuli[17] while provisionally accepting others, suggesting that science exhibits a form of constrained novelty. Whether any published work provokes a change in the order's map is very much dependent on its usefulness to other scientists in furthering their publishable research – and we see this selective mechanism operating in the ongoing critical evaluation of work by other scientists.

Of particular interest is the question of the adaptive response of science when the institutional framework in which it is embedded changes. While we earlier alluded to the effects of central planning on science, probably the more empirically relevant examples emanate from the rise in the 20th century of government funding of science on the one hand and from an increasing tendency of profit-seeking firms organized to undertake fundamental research, as in the 'gen-tech' sector.[18] In terms of our map and model perspective, we would expect that each kind of institutional arrangement would generate different kinds of scientific knowledge by virtue of the allocation of resources associated with different funding sources. Under any given regime of funding, scientists would be expected to orient themselves to research and development activities compatible with the aims of the funding source. Their scientific activities will be receptive to those hunches and conjectures consistent with opportunities for obtaining the resources they require. However, provided the procedures, practices, and conventions that define science (i.e., those institutions consistent with generating reliable and codifiable knowledge) remain operative, each funding arrangement will still produce 'good science', although the particular content of that science is likely to be different. That is, the direction but not necessarily the quality of the outputs generated by science will be affected by different funding and control arrangements, although no doubt over a longer run the attempt to apply central planning to science must be reckoned as self-stultifying and inhibitory to the formation and adaptation of institutions compatible with scientists producing scientific knowledge. This suggests that while it may be possible to understand science in terms of its capacity to produce reliable and codifiable knowledge, such a perspective may be too confining to adequately capture essential aspects of science as an emergent order. In short, looking at science only as a generator of scientific knowledge does not as such tell us very much about the institutional details and mechanisms at play in science that explain its emergent and adaptive characteristics. But the application of the map and model analogy provides a more complete treatment of scientific activity, especially with respect to its capacities and structures for exhibiting anticipatory responses.

The scientific order is divided into various 'communities', defined in terms of domain of inquiry, methodology, ideological commitments, preanalytical visions, and so on. The communities have certain common features that reflect their functioning as anticipatory structures in terms of their receptivity to inputs from individual scientists (both within and external to the community) and their capacity to generate outputs (or scientific knowledge) within and beyond their own confines. Given the low cost of accessing and transmitting scientific knowledge within and especially across communities and the ease with which individual scientists can wear multiple hats as disposition, convenience, and opportunity may allow, we must recognize the nonstatic and even fragile nature of some scientific communities. Despite (or perhaps because of) the possibility of defections, scientific communities as a general rule have developed certain identity-preserving mechanisms geared to the preservation of the community. But these mechanisms, whether we see them as Kuhnian 'invisible colleges' or as indicative of the emergence of a Lakatosian 'hard core', also – and most critically – function as filters that are partially determinative of what sorts of inputs the particular community is likely to be receptive toward. Consequently, these mechanisms cannot be disassociated from the kind of knowledge any particular community is capable of producing or retaining. A community's scientific identity is integral to these filtering mechanisms because that identity, encompassing its stock of accumulated explicit and tacit knowledge (including accepted methodologies and procedures) and the changes that may arise in that stock, is adjudicated by such mechanisms. In the vocabulary of this paper, the content of a scientific community's 'map' includes its identity and its associated filtering mechanisms. If we now imagine a flow of new inputs, say in the form of newly published empirical findings, scientists within the community will form a 'model' of the new environment (i.e., of the relevant domain of inquiry) based on these inputs. But whether this model has the capacity to effect changes in the community's map will be determined by the compatibility between the new model and the existing map as arbitrated by the application of various filtering mechanisms. The essential point for our purposes is that the map and model analogy provides a way to understand the mechanisms at play within science that speak to the adaptive and anticipatory character of the process of the generation of scientific knowledge.

In summary, as we did for the market economy, we characterize the scientific community's functioning as resulting in the generation of a classification – most obviously, a contingent but generally accepted body of scientific knowledge – over a set of inputs. Science, as a social system, is a complexly organized structure with no central locus of control. It functions as an evolving and adaptive order with specialized mechanisms of positive and negative feedback, particularly the publication/citation system, which work to align the overall operation of the order with exogenous environmental changes and constraints and with endogenous changes produced by the system and its components. That science in this form appears to adapt with greater agility and resourcefulness than the available alternatives (such as centrally planned 'science') is, we suggest, precisely attributable to its superior capacity as an adaptive classifying system.

6. Other adaptive classifying systems

Having established a basis for identifying markets and science as instances of adaptive classifying systems, it is not too much of a stretch to think that other social arrangements might also fit the general pattern. Two in particular come to mind:

(1) *City neighborhoods*. Jacobs (1961, pp. 30–40) describes vividly how the basic provision of neighborhood security is a side effect of the interactions of people frequenting the public spaces who, as long-term residents and proprietors, have formed, over time, informal 'networks' of observation and recognition and who stand prepared to warn, assist, or intervene. She is very clear that 'the first thing to understand is that the public peace ... is kept primarily by an intricate, almost unconscious, network of voluntary controls and standards among the people themselves, and enforced by the people themselves' (pp. 31–32). This network, this 'basic supply of activity and eyes', is what provides the necessary background for the support of the influx of strangers vital to the commercial viability of the neighborhood. We can readily see in the arrangements Jacobs describes both a map (the network of residents and shop proprietors and their norms of interaction and intervention) and a model (the ongoing interactions of these people in their various roles and their reactions to perceived security and safety problems).

(2) *Firms*. Our inspiration here is the emerging 'capabilities' theory of the firm, expounded most recently and in most detail by Langlois (1995, 1997) and Langlois & Robertson (1995). These authors focus on the productive activities or skills that people comprising a firm are capable of, and include in these capabilities 'the following activities associated with the profitable production of a good or service: conception, design and development, manufacturing, provision of inputs, marketing and distribution, and many others' (Langlois & Robertson, 1995, p. 11), i.e., routines and interactions that actual people working within actual firms can be observed to engage in. We suspect that most if not all of the insights generated by the capabilities theory would find a comfortable fit within the conception of a firm as an adaptive classifying system. Such a perspective, while of course still having a place for the output-directed activities of the firm, would place emphasis on the structural and operational aspects that link the environmental stimuli impacting the firm to its output responses. It is a commonplace in management literature to point out (without, however, very much in the way of theoretical backing) that a firm's productivity and profitability over the long term is a consequence of its ability, implemented in its internal organization, to learn and adapt. For example, De Geus (1997, pp. 5–9), reporting on an unpublished survey commissioned by Royal Dutch/Shell of 27 large

corporations each at least 80 years old (i.e., with at least double the average life expectancy of Fortune 500-size multinationals), isolates four factors that these firms appear to have in common and that separates them from less successful firms: ability to react quickly to change, a strong sense of corporate identity, a lack of strong central control (especially over experimental activity), and independence in controlling growth and evolution.

None of these factors says anything at all about the products of these firms, except that the firms clearly are consistently able to adapt their product lines to the requirements of their environment (otherwise they would not have survived as long as they have), but all of them hint strongly at the possibility of a characterization of such firms in terms of map and model. It appears that they operate as effective classifiers of the stimuli their employees in touch with their environments can sense (changing customer requirements, current effectiveness of products, and emerging opportunities for expansion, for example), and to do so they must be endowed with effective (and preferably re-entrant) 'connections' between their component employees and departments that enable processes of evaluation and re-evaluation to be applied to the original stimuli and to lead to adaptive reactions. Firms organized in strictly hierarchical style, where transmissions largely flow from management down, are likely to be much less adaptive in general, their ability to adapt to changing circumstances depending to a large extent on the particular abilities of the individual manager in charge and less on the tacit knowledge and interactive opportunities of individual workers to adaptively solve problems.[19] This suggests the map and model perspective might be used in diagnostic, as well as explanatory, mode.

It has also not escaped our attention that some important social arrangements – existing political, legal, and monetary systems, in particular – are less than perfect exemplars of decentralized social systems, and that it may be possible to characterize the perceived failures of such systems in terms of structural impediments to map adaptation or model formation. But whether or not this is a productive line of approach is yet to be seen.

7. Conclusion

Generalizing from Hayek's description of the fundamental structure and operation of the neural order, we have shown that certain social orders (markets, science, and quite possibly others, including city neighborhoods and market firms) can be regarded as operating according to similar abstract principles, although the implementations of these principles differ considerably. 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 (which, in social orders, include effects initiated by the components) 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 (in the form of prices or 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.

Domain	Map	Model	Classification
Neural	Network of interconnected neurons	Current pattern of transmissions within the semi-fixed but mutable neuronal network	The order of sensory qualities as personal knowledge of the environment (which includes the organism's body)
Market	Network of people interconnected via market institutions (property, contracts, and law) and marketing routines and habits	Current pattern of exchange transactions involving transfers of particular goods and services and exchange prices, constrained by the existing market network	The order of market goods and services and market prices as 'market knowledge' of the market's environment (which includes the preferences of market participants)

Science	Network of people interconnected via science institutions (including arrangements of scientific communities) and publishing routines and habits	Current pattern of publication and citation transactions involving published papers, constrained by the existing arrangement of disciplines, journals, and schools (and their protective filters)	The order of reliable and codifiable scientific knowledge organized into disciplines and schools of thought
City neighborhood	Network of people interconnected via observation and conversation in a localized area	Current pattern of alertness to security violations, constrained by the existing possibilities for observation and interaction	The order of basic 'street-level' security as knowledge of the state of the neighborhood
Firm	Network of people interconnected via in-house procedures within an organization geared to producing goods and services	Current pattern of communication involving formal and informal meetings and discussions, constrained by the firm's existing business procedures	The product line (including services) of the firm as knowledge of its market environment

The different implementations of these basic principles in the adaptive classifying systems we have identified are summarized in the following table:[20]

The claim that social arrangements 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, 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 implementations. 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

Earlier versions of this paper were presented at the Behavioral Research Council's conference on 'Dewey, Hayek and Embodied Cognition', Great Barrington, MA, July 18–20, 2003, and at the Southern Economic Association Conference, San Antonio, TX, November 21–23, 2003. We are indebted to Roger Koppl and Bruce

Caldwell for helpful comments on the conference draft, and to the referees and editors of this journal for their careful and insightful advice.

Notes

1. See Hayek (1948, 1979, pp. 199–200, 1994, pp. 138–139).
2. See, for example, Agonito (1975), Weimer (1979, 1982), Miller (1979), Yeager (1984), Herrmann-Pillath (1992), Streit (1993), Butos & Koppl (1993, 1997), Langlois (1997), Smith (1997), Birner (1999), Rizzello & Turvani (2000), and Steele (2002).
3. Hayek was certainly not the first to emphasize the importance for learning of the formation of connections between concurrently active neural elements – see James (1890, p. 566) and Hull (1943, pp. 68–69).
4. Although Hayek's object of interest was the human brain, we note that this is a system that displays certain emergent attributes (consciousness, for instance) that play no part in our generalization. The attributes of neural orders upon which we base our generalization do not include such 'higher' functions. Although consciousness is obviously an attribute of the brains of the people that populate human social systems, there is no reason to ascribe it (or some analog of it) to such systems as a whole. An adaptive classifying system certainly does not have to be conscious to exhibit adaptiveness and even anticipation. In any case, consciousness probably arises as a result of particular structural arrangements that have evolved in certain types of brain, and there is no reason at all to expect that a similar path of evolution would be followed by other, quite different, orders.
5. See, for example, Kauffman (1993) and Holland (1995).
6. Although we are confining our attention to human social systems, we note that (as a referee has pointed out) there exist good examples of adaptive classifying systems elsewhere in the animal kingdom – ant or bee colonies, for example. See Landa & Tullock (2003) for an interesting look at such systems.
7. The term 'classifier systems' has been used in the artificial intelligence field in the context of work on 'learning classifier systems' (Holland 1980; Holland et al. 2000), and it is interesting to consider, as an aside, how closely these systems conform to our elaboration of the essential attributes of adaptive classifying systems. In a learning classifier system, the components are production rules (consisting of a tag or condition part and an action part) called 'classifiers'. A stimulus from the environment is a message that is matched against the condition parts of the population of rules, and matching rules can become active, with their action parts initiating new messages or contributing to motor responses. Which matching rules are activated can depend on their 'strength', a parameter which is increased if a 'payoff' is received from the environment during the time they are active, and this reinforcement can be distributed back to rules responsible for messages that resulted in the activation of the rewarded rule. New rules are produced by the application of a genetic algorithm to existing rules, with rule strength serving as a measure of fitness for parenthood; rules of low strength are candidates for elimination. We can identify the map in this arrangement as the collection of production rules, and, to the extent that one rule may invoke other rules indirectly through the posting of messages, there is some level of connectivity between the components. The model is implemented through the posting of messages and the resulting activation of matching rules. The map is modified, not directly as a result of activity, but by the effects of 'payoff' feedback and, even more indirectly, as a result of the production of combinations of high-strength rules and the elimination of low-strength ones. Whether or not the loose interconnectivity of components and the indirect and somewhat *ad hoc* mechanisms for map change might constitute weaknesses of the system, we cannot definitively say. But, although it was not our intention to extend our work into the domain of artificial systems, it is heartening to see that there is at least some basis for thinking that our generalizations might be applicable there, too. We do note, however, a significant difference in organization – in learning classifier systems, the adaptive capability

arises because the components are subject to a form of selection, whereas we are talking about systems in which the component interactions result, as a side effect, in a form of structural self-organization that is the basis for the adaptive capability.

8. Hayek (1967, pp. 73–74), in the essay entitled 'Notes on the Evolution of Systems of Rules of Conduct'.
9. The obvious differences between people and neurons – that people have goals and are motivated to pursue them, and that people are themselves knowledge-generating objects – do not limit the applicability of our analogy as much as one might expect. We treat the unique goals and innovative contributions as localized inputs to the system to which the system will adjust. To put it succinctly, people are taken to be neuron-like in their formation of mutable interactions with each other, sensory-organ-like in their contribution of stimuli that activate and alter the system, and motor-organ-like in their ability to effect changes on the system's environment.
10. This is the kernel of truth in Galbraith's (1958) 'dependence effect' – but see Hayek (1961) for an antidote to the further conclusions Galbraith draws from it. See McQuade (2000) for explicit simulation of the effect on market clearing in an agent-based model.
11. The various sorts of feedback mechanisms characteristic of the operation of complex social systems suggests the difficulty of treating any of these factors, whether exogenous or endogenous, as 'originating' in any causal sense of that term. The fact that the external environment has been systematically and irrevocably altered by human societies or that an agent's novel ideas are built upon an inherited stock of social knowledge argues against any straightforward assignment of truly causal factors, whether external or internal to the system. We can side-step such difficulties because our interest here is the more modest one of viewing how stimuli affect the system in the context of the model and subsequently the map. Consequently, this allows us, with the aforementioned qualifications, to assume some starting point of the system in analyzing its response to stimuli of various kinds and the salient kinds of feedback effects such stimuli induce.
12. The story here is complicated by the fact that an agent's behavior may be influenced by his expectation that another agent may act in a particular way.
13. It might seem from our description that the role of entrepreneur is missing, but this is not so. The activities of entrepreneurs, reacting to their particular perceptions of their environment, provide inputs to the market system that induce follow-on activity in the system's model and may end up inducing changes in the system's map. This is not fully captured by characterizing such changes as 'bringing the system back toward equilibrium' – it is rather more complicated than that, as we shall describe below. Moreover, modeling agents such as cognitively heterogeneous is quite compatible with our treatment and provides an entry point for treating agents as creative, innovating, and alert entrepreneurs.
14. The functional distinctiveness of the market order resides precisely in these capabilities, that is, in its adaptive properties, and it is such properties that allow us to systematically address the relative merits or failings of alternative economic arrangements. Such considerations ought to figure prominently in the scientific analysis of different institutional frameworks.
15. We think (see McQuade & Butos 2003) that modeling scientists as self-interested reputation-seekers is an analytically and descriptively useful way to understand the coordinating mechanisms of science.
16. With the advent of the Royal Society and its journal, the *Philosophical Transactions* – see Merton (1973, pp. 191–203) and McQuade & Butos (2003).
17. For example, scientists who previously engaged in fraudulent research are likely to have future research discounted despite its inherent quality. In science, reputation seems to rule.
18. Throughout most of the 1990s in the US the Federal Government provided about 30% of total R&D funding, private sector industry about 65%, and other institutions (state governments, universities, and non-profits) about 5%. The growth rate of non-federal funding over the past decade has been about three times that of the Federal Government (Payson, 1999). Since 1997, non-federal R&D spending has averaged about 73% of the total, while that of the Federal Government about 27% (NSF, 2003).

19. The role of tacit knowledge and capabilities in expanding the adaptive qualities of organizations is an unexplored dimension in analyzing the potential for successful corporate mergers and acquisitions.
20. Such a table, inevitably, oversimplifies the issue, tending to conflate morphology and function. The 'classification' – the particular kind of knowledge generated – of each domain is not separable from the map and the model; in fact, the map and the model constitute the specific functionalities of how we account for the emergence of the respective classifications. Even for the neural domain, there is no clear separation between the processes by which knowledge is generated and the output of (the classifications produced by) the order.

References cited

Agonito, Rosemary. 1975. Hayek revisited: mind as a process of classification. Behaviorism 3:162–171.

Birner, Jack. 1999. The surprising place of psychology in the work of F.A. Hayek. History of Economic Ideas 7(1/2):43–84.

Boehm, Stephan. 1994. Hayek and knowledge: some question marks. Pp. 160–177 in M. Colonna, H. Hagemann & O. F. Hamouda (ed.) Capitalism, Socialism and Knowledge: The Economics of F.A. Hayek, vol. II. Edward Elgar, Aldershot.

Buchanan, James M. 1979. What should economists do? Liberty Press, Indianapolis.

Buchanan, James M. & Viktor Vanberg. 1991. The market as a creative process. Economics and Philosophy 7(2):167–186.

Butos, William N. & Roger Koppl. 1993. Hayekian expectations: theory and empirical applications. Constitutional Political Economy 4(3):303–329.

Butos, William N. & Roger Koppl. 1997. The varieties of subjectivism: Keynes and Hayek on expectations. History of Political Economy 29(2):327–359.

Butos, William N. & Thomas J. McQuade. 2002. Mind, market, and institutions: the knowledge problem in Hayek's thought. Pp. 113–133 in J. Birner, P. Garrouste & T. Aimar (ed.) F.A. Hayek as a Political Economist. Routledge, London.

De Geus, Arie. 1997. The living company. Harvard Business School Press, Boston.

Edelman, Gerald M. 1987. Neural Darwinism. Basic Books, New York.

Edelman, Gerald M. 1992. Bright air, brilliant fire. Basic Books, New York.

Galbraith, John K. 1958. The affluent society. Houghton Mifflin, Boston.

Hayek, Friedrich A. 1948. Letter of 6 Nov. to J. Nef in the Hayek archives, box 55, folder 1. Hoover Institution, Stanford.

Hayek, Friedrich A. 1952. The sensory order. University of Chicago Press, Chicago.

Hayek, Friedrich A. 1961. The *non sequitur* of the 'dependence effect'. Southern Economic Journal (April):346–348.

Hayek, Friedrich A. 1967. Studies in philosophy, politics and economics. University of Chicago Press, Chicago.

Hayek, Friedrich A. 1973. Law, legislation and liberty, vol. 1: Rules and order. University of Chicago Press, Chicago.

Hayek, Friedrich A. 1975. Kinds of order in society. Institute for Humane Studies, Menlo Park.

Hayek, Friedrich A. 1976. Law, legislation and liberty, vol. 2: The mirage of social justice. University of Chicago Press, Chicago.

Hayek, Friedrich A. 1979. Law, legislation and liberty, vol. 3: The political order of a free people. University of Chicago Press, Chicago.

Hayek, Friedrich A. 1994. Hayek on Hayek: an autobiographical dialogue. S. Kresge & L. Wenar (ed.) University of Chicago Press, Chicago.

Hebb, Donald O. 1949. The organization of behavior. Wiley, New York.

Herrmann-Pillath, Carsten. 1992. The brain, its sensory order and the evolutionary concept of mind: on Hayek's contribution to evolutionary epistemology. Journal of Social and Biological Structures 15(2):147–187.

Holland, John H. 1980. Adaptive algorithms for discovering and using general patterns in growing knowledge bases. International Journal of Policy Analysis and Information Systems 4(3):245–268.

Holland, John H. 1995. Hidden order. Addison-Wesley, Reading, MA.

Holland, John H., Lashon B. Booker, Marco Colombetti, Marco Dorigo, David E. Goldberg, Stephanie Forrest, Rick L. Riolo, Robert E. Smith, Pier Luca Lanzi, Wolfgang Stolzmann & Stewart W. Wilson. 2000. What is a learning classifier system? Learning classifier systems: from foundations to applications. Springer-Verlag Lecture Notes in Artificial Intelligence 1813:3–32.

Hull, Clark L. 1943. Principles of behavior. Appleton-Century-Crofts, New York.

Hull, David L. 1988. Science as a process: an evolutionary account of the social and conceptual development of science. University of Chicago Press, Chicago.

Jacobs, Jane. 1961. The death and life of great American cities. Vintage Books, New York, 1992.

James, William. 1890. Principles of psychology, vol. 1. Dover, New York, 1950.

Kauffman, Stuart A. 1993. The origins of order. Oxford University Press, New York.

Landa, Janet T. & Gordon Tullock. 2003. Why ants do but bees do not construct satellite nests. Journal of Bioeconomics 5(2–3):151–164.

Langlois, Richard N. 1995. Do firms plan? Constitutional Political Economy 6:247–261.

Langlois, Richard N. 1997. Cognition and capabilities: opportunities seized and missed in the history of the computer industry. Pp. 71–94 in R. Garud, P. Nayyar & Z. Shapira (ed.) Technological Entrepreneurship: Oversights and Foresights. Cambridge University Press, New York.

Langlois, Richard N. & Paul L. Robertson. 1995. Firms, markets and economic change. Routledge, New York.

McQuade, Thomas J. 2000. Orders and institutions as objects of investigation. Ph.D. dissertation, Auburn University.

McQuade, Thomas J. & William N. Butos. 2003. Order-dependent knowledge and the economics of science. Review of Austrian Economics 16(2–3):133–152.

Merton, Robert K. 1973. The sociology of science. University of Chicago Press, Chicago.

Miller, Eugene F. 1979. The cognitive basis of Hayek's political thought. Pp. 243–267 in R. L. Cunningham (ed.) Liberty and the Rule of Law. Texas A & M University Press, College Station.

Mises, Ludwig. 1966. Human Action. Henry Regnery, Chicago, 3/e.

NSF. 2003. National patterns of research and development resources: 2002 data update. National Science Foundation 03-313. http://www.nsf.gov/sbe/srs/dst.htm.

Payson, Steven. 1999. National patterns of R & D resources: 1998. NSF 99-335. National Science Foundation, Arlington.

Rizzello, Salvatore & Margherita Turvani. 2000. Institutions meet mind: the way out of an impasse. Constitutional Political Economy 11(2):165–180.

Smith, Barry. 1997. The connectionist mind: a study of Hayekian psychology. Pp. 9–30 in S. F. Frowen (ed.) Hayek: Economist and Social Philosopher. Macmillan, London.

Steele, Gerald R. 2002. Hayek's sensory order. Theory and Psychology 12(3):125–147.

Streit, Manfred E. 1993. Cognition, competition, and catallaxy: in memory of F. A. Hayek. Constitutional Political Economy 4(2): 223–262.

Weimer, Walter B. 1979. Notes on the methodology of scientific research. Lawrence Erblaum, Hillsdale.

Weimer, Walter B. 1982. Hayek's approach to the problems of complex phenomena: an introduction to the psychology of *The sensory order*. Pp. 241–285 in W. B. Weimer & D. S. Palermo (ed.) Cognition and the Symbolic Processes, vol. II. Lawrence Erblaum, Hillsdale.

Yeager, Leland B. 1984. Hayek and the psychology of socialism and freedom. The AEI Economist (November):1–5.