

A SYSTEMS SCIENCE FRAMEWORK FOR UNDERSTANDING THE NATURE OF GOVERNANCE

George Mobus

Author Address

ABSTRACT

The natural world of life is replete with examples of systemic governance subsystems that operate to sustain the continuance of those systems. Every cell, organism, population, and ecosystem demonstrates various self-regulation and environmental coordination mechanisms that have evolved to ensure the long-term viability of that system. A formal approach from systems science that is built on these natural governance subsystems may provide some guidance to our understanding of human social systems and their governance. The emergence of higher levels of organization in the origins and evolution of life can be seen to be the story of increasing sophistication in governance subsystems as disparate complex adaptive systems coalesce into “societies” of interacting entities (super-molecules to primitive protocells, prokaryotic cells to eukaryotic cells, those to multicellular organisms, those to communities, etc.). At each stage in this on-going emergence of higher levels of organization the one consistent aspect is how hierarchical cybernetic structures have contributed to the stabilization of functional relations among the component entities leading to sustainable super-entity structures. The progression is from simple cooperation of multiple entities to intentional coordination emerging to manage complexity. Information processing and decision subsystems (agents) that took responsibility for logistical coordination among components and others that managed tactical coordination of the whole system with external (environmental) entities, resources, and threats evolved to keep increasingly complex biological entities able to maintain their existence and reproduction. Now the governance of human social systems that seek to exist in some kind of harmony with the Earth’s ecology (what I call the Ecos) has emerged in the last 100k years or so and evolved over that time frame to produce the modern socio-economic systems in existence today. But it (characterized here as the neoliberal capitalistic democracy) is not as evolved as, say, the mechanisms of metabolic regulation. There are numerous reasons to believe that the modern governance subsystem is, in fact, moving human societies toward the opposite of sustainable existence. A systems examination of the theory of governance subsystems (hierarchical cybernetics) suggests pathways toward a more functional governance subsystem for human societies. The theory covers the regulation of economic flows as well as the legal superstructure and moral/ethical aspects of culture that collectively constitutes the governance subsystem of a human society embedded in a meta-system, the Ecos.

Keywords: Governance, Systems Science, Hierarchical Cybernetics, Social Emergence, Decision Agents

INTRODUCTION

We are faced with an extremely difficult set of problems as we acknowledge the Anthropocene as a new epoch. In fact one could easily argue that this age of geological markers put down by human activity is the result of a failure of our species to govern our interactions with the world and with each other. Our species is effectively out of control in terms of living in balance with the Earth as it has existed since after the Cretaceous–Paleogene (K–Pg) extinction event. As a result our activities are altering the atmosphere, the hydrosphere, the biosphere, and even the lithosphere. Nature is largely based on the idea that the complex cycles of those spheres interact with both positive and negative feedbacks that keep the environment on the surface relatively friendly to life in general (Lovelock, 2006). In some sense the Earth ecosystem (which I refer to as the “Ecos”) has achieved a relative dynamic steady state condition not unlike the pre-K-PG boundary. The biosphere has some resemblance to a mature body, which has developed and grown to its maximum size, or a quasi-stable climax ecosystem. Humans, through higher-order cognitive capacities, have broken out of the steady state and now almost resemble a kind of cancerous growth that threatens the rest of the biosphere, or at least a significant portion of it.

It is ironic that this phenomenon is linked to the human desire to thrive and an abundance of intelligence for creating tools and technologies to presumably support that desire. Too much of a good thing is turning out to be not very good after all.

It is also ironic that just as we are beginning to feel the impacts of the rapid changes that underlay the Anthropocene that same level of intelligence has allowed us to become conscious of what is happening through our invention and use of science. I think this is profoundly important. I propose that, in particular, the science of systems (systems thinking, systems approach, etc.) is a body of knowledge and a way of conceptualizing the world that holds a key to both understanding the whole of what is happening and to seeking a holistic set of solutions that might mitigate the worst effects of what the Anthropocene might offer.

The State of Systems Science

Sometime during World War II, in the United States and parts of Europe several threads of scientific investigation began to coalesce owing to their subjects having very strong relations. The war itself was a catalyst to promoting and bringing fields such as communication and information theory, control theory, computation theory, and others together under a loose rubric we recognize as systems science. At the same time, and especially shortly after the war, scientists from many different fields were coming to appreciate the idea of a unity of concepts that underlay all of the sciences, natural and social. During the forties and fifties there were many scientists who saw the grand unification of these various threads and began promoting the universal patterns of how things worked – general systems theory (von Bertalanffy, 1968).

But over the next five to six decades rapid progress within each of the original sub-fields, which meant pursuing traditional reductionist approaches, and coupled with a new

A Systems Science Framework for Governance

emphasis in the tenure process at major universities on research publications in narrow domains, proceeded to create the same kinds of silos seen in ordinary sciences. Rather than integrating the concepts from these fields, they tended to develop in parallel, knowing something about each other but still focused on their particular viewpoint. Today we have network science, complexity science, cybernetics, information theory (and science), operations research, system dynamics, and so on. Interestingly writers in each of these fields refer to systems in the general sense, but tend to subsume the general concepts of systems under their own field's title.

Of course some kind of attention to details using reductionist methods is necessary but it is not sufficient to really grasp the whole nature of systems. Today it should be possible to think again of the reunification of these various areas of knowledge as they address a full understanding of how things work (Mobus and Kalton, 2014). This paper attempts such an approach as it applies to the concern for governance of human society, especially in the Anthropocene.

Old Ideas with a New Framework

A number of old ideas will be presented here, but what is new is a framework that builds from a reintegrated vision of systems science.

This paper seeks to introduce the idea that a whole (and general) systems approach to governance is a potential first step toward moderating the Anthropocene. The fundamental idea is that nature is replete with examples of evolved governance systems for a variety of living systems and meta-systems (c.f. Beer, 1980, 1981). If these are understood they might provide models that could be used to design governance structures and mechanisms that will moderate the unbridled human craving for comfort and convenience that is largely responsible for our current state of affairs. The pathway to such understanding begins with a framework for analysis. There are a set of "principles" of systems science that have been distilled from the works of so many systems thinkers and theorists (Mobus and Kalton, 2014). Using those principles I propose an outline for that framework. The new ideas to be introduced are:

- Complex Adaptive and *Evolvable* Systems (CAES)
- Societies become entities in nature (Emergence), Bourke, 2011
 - Major transitions, Smith & Szathmáry, 1995
 - Hierarchy of emergences, Morowitz, 2004
 - How evolution solves the governance problem in nature
- Hierarchical cybernetic model (not just simple homeostatic feedback)
- The problem with human decision makers – lack of adequate wisdom

A Systems Science Framework for Governance

This is only an outline. I am developing a series of books that will add some flesh to this mere skeleton. The first will deal with the last bullet item, the condition of human mentality that leads to often faulty decisions (in terms of short-term and local scale thinking) with respect to the scale and scope of the ("wicked") problems we need to solve. The basic problem turns out to be that we humans have not evolved a sufficient capacity for species-level wisdom. This is a function of our brain development and I call this capacity "sapience" to recognize that fact (Mobus, 2015 in development).

Governance, the subject of a second book in this series (Mobus, 2016, in early development), whether found in cells or corporations, depends on effective decision agents that are the proximate cause of corrective actions (which can be as simple as a thermostat or as complex as formulating policies). Human decision makers suffer from a number of deficits that result in faulty decision making; a prime reason why organizational governance systems frequently falter or fail. This will always need to be taken into account when designing a governance system to bring our species into accord with the rest of the Ecos.

The answers (if they exist) will be found in the nature of complex systems and our understanding of them as phenomena.

One category of complex systems that has been gaining increasing interest of late is the complex adaptive and evolvable system (CAES). A number of authors conflate the ideas of adaptivity and evolvability and leave the term complex adaptive systems (CAS) to cover both concepts. This confusion comes from the characterization of evolution as being species-level adaptation without explicitly saying so. In this author's view that fails to recognize some fine distinctions that should be considered in understanding the very long term dynamics of these systems. For example biological individual entities prior to more advanced neocortical brains (e.g. mammals and birds) have different degrees of adaptivity to variations in their environmental conditions but they are not evolvable since they do not create completely new behaviours (or their underlying mechanisms). More advanced mammals, and especially humans, can learn new behaviours due to evolutionary-like learning in their neocortices. In that sense humans are evolvable systems. Human social systems are clearly evolvable. As more is being understood from the biological examples of evolvability we are beginning to recognize the differences between the two kinds of adaptivity. Social systems (to be defined below) are CAESs embedded in very complex, uncertain, and non-stationary environments that exert long-term selective forces on those systems. Long-term sustainability depends on a particular subsystem function of CAESs that we derive from hierarchical cybernetics - governance. However, another, more subtle difference between individual, or entity, adaptivity and evolvability is that the former is found in CASs such as causally closed bounded systems as are the subject of system dynamics studies. CAESs are causally open, meaning they can alter internal structures to acquire new functions or eliminate old ones. Very little is known about such systems because currently there is no language for expressing their capabilities that combine system dynamics and evolvability. In part, this paper should point to possible developments for such a language. Governance is intrinsic within a bounded but "open" system and should include provisions for evolvability (e.g. the US Constitution provides for the creation of amendments as needed).

A Systems Science Framework for Governance

In this paper I explore the aspect of CAESs that gives them certain qualities relative to long-term persistence and sustainability of core functions. I set out a framework of such exploration based on a set of principles from systems science (Mobus & Kalton, 2014). In particular I will bring various threads of research together in examining the governance of CAESs found in nature that suggest something like a general governance theory that might provide insights into what the human social system governance could evolve into, if, indeed, human social systems are evolvable. The analysis provides evidence that human social governance, in its current forms, is still immature as compared with more evolved systems in other domains of the natural world.

The viewpoint taken here is that of the "major transitions" or "hierarchical emergences" views of researchers in prebiological and evolutionary processes (Bourke, 2011; Morowitz, 2004; Smith and Szathmáry, 1995). A major feature of this point of view is the predominant role of cooperation and coordination among systems of subsystems that comprise emergent entities. Examples include cooperation and coevolution of nucleotides and proteins (and other prebiological molecules) leading to the first cells, endosymbiosis leading to the first eukaryotes, and numerous examples of symbiotic relations leading to multicellular systems. All of these have been described as societies of disparate systems and those societies evolved to form distinct entities at a higher level of organization (Mobus and Kalton, 2014, chapter 10). Biological evolution has been the story of increasing levels of organization, producing most recently societies of mammals. A few of these are described as 'eusocial' (Wilson, 2013). Humans are among them. Thus taking the collectives of humans that we call societies as (at least) potential entities, in the major transitions sense, we can examine the emergence of mechanisms that enhance cooperation and coordination.

Though the concept of eusociality (or what some have labelled "hyper-sociality") is cogent to the viewpoint taken here, the subject is beyond the scope of this paper. A fuller account will be available in Mobus (2015, in preparation). For our purposes we will claim without proof that eusociality in humans accounts for the human tendency to cooperate, thus forming a substrate for mechanisms such as markets and group undertakings for mutual benefits. Sober & Wilson (1998) have provided the arguments for how eusociality emerges in human societies through the evolutionary process of group selection. I will adopt those arguments here.

The central question I seek to answer is: How did entities emerge and evolve capabilities to self-regulate and achieve stability, resilience, robustness, adaptivity, and evolvability leading to their very long-term continuance through geological time scales? The question is really, how did living systems persist to the present as autonomous entities? The answer lies in the emergence and development of hierarchical cybernetic subsystems that provided the kind of "management" framework that led to those qualities listed above.

COMPLEX ADAPTIVE AND EVOLVABLE SYSTEMS

Complex adaptive systems (CAS) can be found throughout biology and even some manmade entities. Evolvable systems are a little more difficult to characterize. In this

A Systems Science Framework for Governance

section I want to clarify the difference between adaptivity and evolvability insofar as they operate in specific kinds of systems. The two terms represent differences in the underlying mechanisms that allow a system to accommodate changes in its environmental conditions and over different time scales.

Adaptivity vs. Evolvability

An adaptive system (entity) is capable of modifying its short-term behaviour in order to accommodate some variation in environmental conditions that place the entity under stress. Examination of the internal mechanisms within such entities reveals they are flexible in ability to operate nominally over a range of variations that, as is the case, are representative of what can happen in that environment under normal conditions. In biological entities this can involve responses to varying stimuli that are attempts to get the entity into conformance with the environment and still meet its existential needs. Homeostasis, physiological adaptivity, etc. are examples. When muscle mass increases in response to weight lifting the individual's body is adapting to the new demands placed on it. The mechanisms for adapting are already in the design of the entity and merely need to be "exercised."

Adaptation in an individual is generally limited to relatively narrow ranges of critical parameters such as temperature, energy, oxygen, and water availability, etc. Their internal mechanisms can operate more or less optimally as long as those parameters remain within the nominal ranges. If they get too far out of these ranges then the entity suffers stresses that can drain the system and possibly cause damage. The individual cannot evolve a new capability to live in a very different environment.

Evolution is the process of adaptation of a species (the population of entities taken as the system of interest). The underlying mechanism depends on an ability to generate novelty in terms of functional capabilities among some members of the population. This can only succeed in a population context; there have to be numerous copies of the same basic system in which a few variations in functionality can be "tested." The environment provides the test platform. Some variations will prove to be "better" (more fit) in a given environment and tend to be replicated more frequently than the old normal model. This, of course, is a description of neo-Darwinian biological evolution but it also applies to cultural evolution. Some of the more interesting questions for humanity have to do with the interplay between the biological evolution of humans and the evolution of culture, called coevolution. This is a large part of what gives rise to societies as we have witnessed them.

Learning in the Neocortex

Learned behaviours involve modifications in behavioural sequences that already exist within the capacity of the brains of animals. In other words, the ordinary sort of learning in most animals (pre-neocortex), e.g. conditioned responses in invertebrates (Alkon, 1987), is an advanced form of adaptive behaviour. In neocortical brains a mechanism that resembles evolutionary process dominates the learning. Mammals can learn new behaviours that are not clearly of the conditioned response type. Such behaviours are new

A Systems Science Framework for Governance

concatenations of lower level or atomic behaviours but the ability to make concatenations that are novel and can then be tested by environmental contingencies starts to fit the evolvability model to some degree.

The human neocortex expands on the basic learning in mammalian neocortices, especially in the prefrontal cortex. Humans learn concepts that actually resemble populations of species (Mobus and Kalton, 2014, section 8.25), i.e. multiple copies of similar concepts that can be altered (sub-concepts added or subtracted) and then tested as the world evolves. In this sense human learning includes an evolvable capability that makes our species unique among individual biological entities. We can learn about a possible future (think about the future) and test what we learn by observing what actually develops. Thus, humans represent both an adaptive and an evolvable system. It is no surprise then that societies of humans are also CAESs.

Autonomy

CAESs are also autonomous entities. They are capable of not only seeking goals but of setting them as well. In the biological and supra-biological world (e.g. ecosystems and organizations) goals may be complex and even fuzzy, but in all cases they must ultimately support the biologically mandated goals of maintaining life, obtaining necessary resources, sustaining activity, and growth. The governance of autonomous CAESs includes mechanisms that align higher-order “invented” goals with these biologically mandated ones. The relation between autonomy and strategic management will be made explicit below.

Modelling CAESs

There are a number of approaches to modelling systems in general and dynamic systems in particular (c.f. Meadows, 2008). The problem with, for example, typical system dynamics modelling has been that only static structural systems (i.e. the input/output, flows and stocks, and controls) are permitted. The system is compiled in computer code that, with a little extra work, might be able to represent adaptability, but so far as I am aware, does not provide any built in mechanisms for evolvability. The only way to evolve a system dynamics model is to add on additional model elements by hand and then recompile the system to see what happens. There is no endogenous change in functions during runtime.

It is beyond the scope of this paper, but the possibility of including evolvability into a system dynamics-like modelling language is being investigated. This approach goes beyond the typical evolutionary or genetic programming paradigms by including specific mechanisms for generating novel functionality (and structures) by copying existing ones and then applying stochastic modifications to the copied version to see if a novel function that improves overall fitness emerges. This model of emergent complexity is seen in natural systems. The copying process is activated by creating selection pressures that push the limits of adaptability or introduce new exogenous entities, such as a new resource that the system might be able to take advantage of. The whole system must operate in a framework that tests its fitness over time. The output of such a modelling language would be not only graphs of state variables (stocks) over time, but also include

A Systems Science Framework for Governance

the non-linear effects of evolved functionality. This work is just getting underway but looks promising.

SOCIETIES IN EVOLUTION

In recent decades more evolutionists are recognizing the important role of cooperation between disparate entities that form stable and strongly coupled interactions. Bourke (2011) calls such super-entities “societies” and describes how such groups of individual entities begin to become so dependent on one another that they form an effective entity at a higher level of organization. The boundaries are often defined by a change in coupling strength between member subsystems that are in the entity versus entities that are outside. The interactions between some members of the society and the environment might be of varying strength but are clearly sparse compared with interactions internal to the entity.

Societies: From Aggregates to Collectives

The progression of grand-scale evolution begins with the aggregation of component entities, themselves systems capable of some degree of adaptation. They are able to interact weakly with other entities and can form loose affiliations. This is a form of self-organization as it is not particularly guided by external forces but neither is it disrupted by any forces. Such an aggregation is able to obtain a function in the sense that it may obtain some resources from other entities as well as produce products that are absorbed by other entities (see Figure 1 below). If that arrangement is beneficial to the receiving entity lines of communications may provide feedback that causes the forming society to strengthen its internal structure and its boundary conditions. The society has become a precursor entity capable of behaving as a whole and further defining itself from its environment.

If the new entity is also evolvable then it can add new functions (new subsystems) by various means. Such new subsystems are at first novelties, but should they acquire independent and helpful functions that increase the overall fitness of the entity then they become new entities, essentially. All of that happens in the context of there being a population of similar entities where individuals can explore the space of possibilities independently.

A Systems Science Framework for Governance

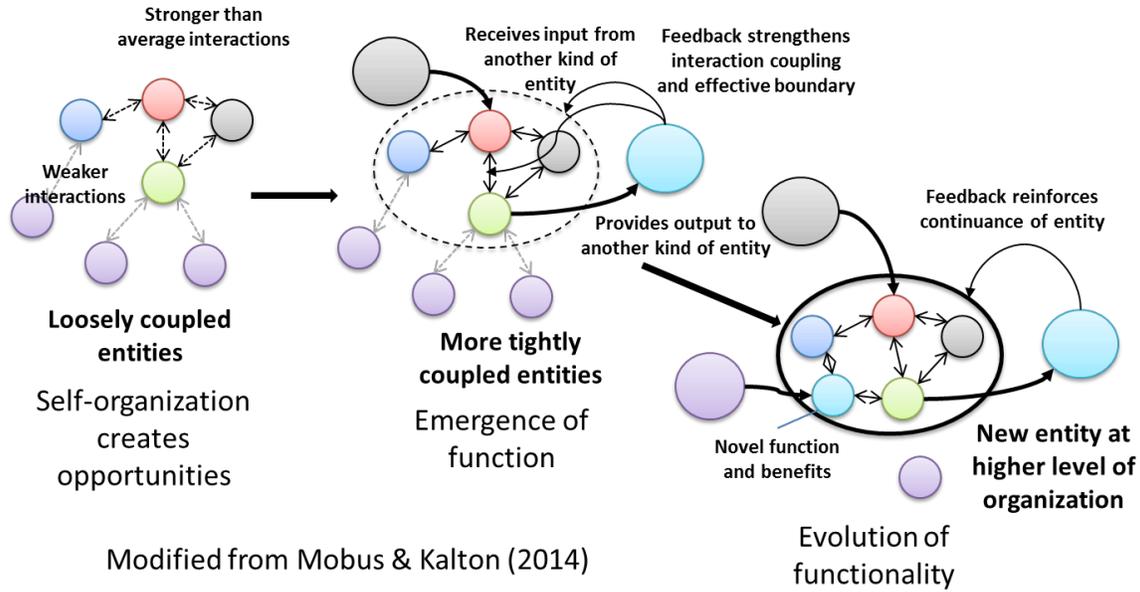


Figure 1. Societies are precursors of entities (whole systems) at a higher level of organization.

Note that a new level of organization has emerged. New entity types with new kinds of interactions and functionalities have developed that may have surprising new properties, individually and collectively. The cycle can continue so long as there is a sustaining flow of free energy with which to construct bonding interactions (such as moving products from producers to consumers).

The Emergence and Evolution of Cooperation

The details of how the coupling of interactions increases are of prime interest. They will, of course, vary depending on the levels of organization at which we find them. Chemical bonds forming in the pre-biotic domain are quite different mechanisms from the lock-and-key mechanisms found in high molecular weight macromolecules such as enzymes and their substrates. Still more different are the mechanisms of oxytocin-induced bonding in human groups! Yet the dynamics of progression under evolutionary processes is the same.

In human beings the evolution of cooperation is at the root of forming societies of individuals and then observing how those societies begin to function as precursor entities. We even assign proper nouns to these precursors, giving names to clubs, churches, companies, towns, and nations. The behavioural repertoire that humans engage to form cooperative efforts is generally understood as some form of eusociality. There are now understood to be a number of mechanisms at work in the brain to promote cooperation (Mobus, 2015, in preparation). For example, empathy is a capability that attunes an individual to the emotional states of others, thus forming one kind of strong

A Systems Science Framework for Governance

communicative interaction. In normal humans empathy results in individuals “feeling” for the other and that sets one’s mind to want to behave in a way helpful to the other.

Another and extremely powerful interaction is through language. A considerable literature on the structure and functions of language has been produced so I will not attempt to encapsulate it here. The reader is especially encouraged to investigate the works of Tomasello (2014) regarding language as a mechanism (among many) that promotes cooperation.

The clear message from many examples of cooperation in different forms at different levels of organization (in the emergences viewpoint) is that it provides considerable benefits for the co-operators in terms of creating stable formations that tend to derive higher fitness than can be achieved by individuals alone.

Human beings are not eusocial in the same way ants or bees are. We are not cooperative because we are genetically determined to behave as automata whose instinctive actions are triggered by pheromone signals (though there is some of this in our biology). Moreover, our mechanisms for cooperativity seem to be easily overridden or swamped by cultural and situational circumstances (more later). Indeed there is some evidence that suggests we are really just nascent entities with respect to the evolution of those mechanisms. We do form societies (broadly defined), but they are under constant flux due to cultural evolution and subsequent selection forces. We evolved to be eusocial, but only just. In our more primitive tribal forms natural governance structure and function emerged in those simple societies to support the form of evolution based on group selection. But our immense success as a species has led to many orders of magnitude expansion of the scales of societies. And that has led to levels of complexity that far transcend the biologically-based forms of cooperation. What has emerged now is a set of mechanisms that reinforce and/or support cooperativity in one sense, but also admits a role for intense competition that may be proving the undoing of all we achieved by being cooperative.

Evolution of Complexity and the Need for Governance

A natural definition of complexity (one we routinely see in nature) is that it involves numbers of kinds of atomic entities (heterogeneity), numbers of each of those kinds (populations), numbers of interactions between kinds (network densities), and levels of organization (in the Simon sense). Such a definition does not preclude other aspects of complexity such as the fact that interactions can be nonlinear or lead to chaos or that the levels might not have some obvious fractal qualities. But those have to be incorporated into a framework in which complexity means big, complicated, and hard to parse. This is generally what we mean by complexity in living systems and it is certainly the case for CAESs.

As we saw above, complex systems evolve from the interactions between simpler systems coming to be obligate; societies that start out loosely defined evolve into entities at a higher level of organization. From the perspective of the whole, the subsystems come to rely on one another for material and energy flows. But with complexity comes the

A Systems Science Framework for Governance

central problem of subsystems succeeding in working together. Cooperation (e.g. mutualism) cannot alone ensure the long-term stability of the relations and the continuance of the new super-system. Corning (2010, p-116) asks, “Can hierarchical, cybernetic controls evolve spontaneously...?” If he means by spontaneously something like instantaneously *de novo* then it should be clear that is not the case. But cybernetic capabilities can emerge from existing cooperative relations by normal evolutionary processes.

Along with the emergence of beneficial interactions we note the emergence of information channels that allow subsystems to communicate needs and intentions to one another in order to facilitate the flows of materials and energies that constitute the production systems (Figure 2 below). Later some subsystems come to specialize in assisting the coordination of other productive subsystems. Their “reward” for providing this computational service is energy (and repair services as needed).

A first step in this direction is the development of special computational processes that mediate feedback control to a main work process entity (Figure 3 below). Figure 3 shows an advanced form of such an information gathering and processing entity. This is the basic feedback control mechanism found in homeostatic mechanisms in nature and human management (management by exception) processes. In this version the entity has evolved a more sophisticated decision processor (computer, analogue or discrete), a stored model of the mapping of inputs to outputs (the control model), and a means for sensing the measured parameter and comparing it to an ideal value (set point), generating an error signal which is the technical information used by the system to “make” a control decision.

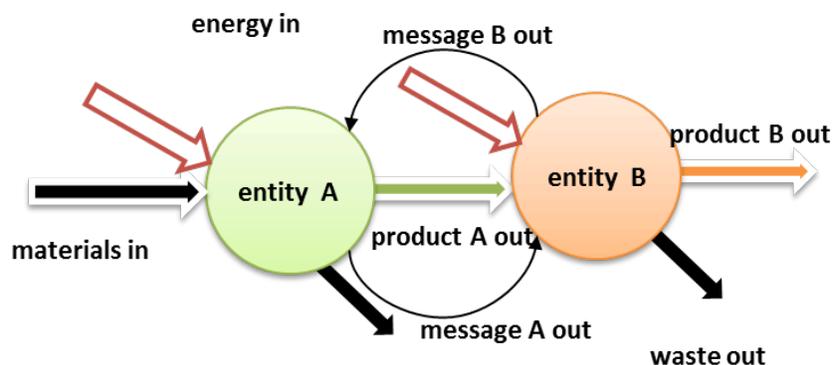


Figure 2. Two product-coupled entities evolve communications channels that allow them to cooperate and match the flows of product A to the needs of entity B.

A Systems Science Framework for Governance

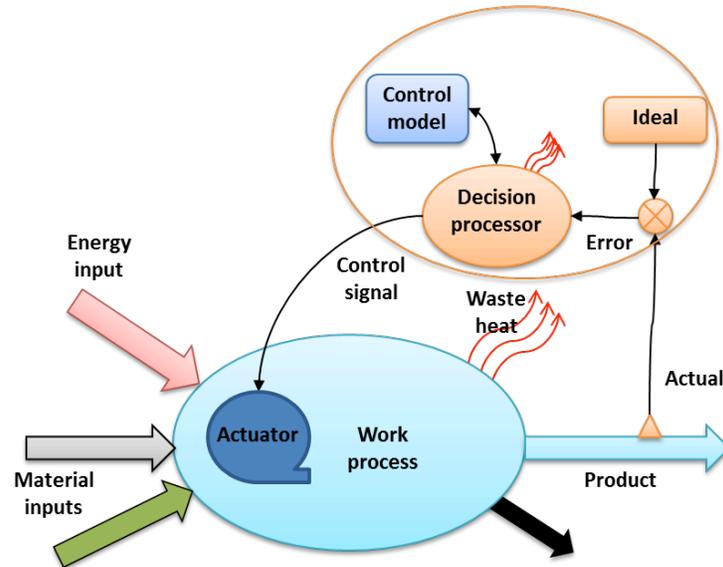


Figure 3. A process that previously may have received the flow of product from a work process (e.g. in Figure x above) evolves to use less product and only measure some quality aspect of the product (e.g. flow rate) and provide control feedback to actuators in the supplier entity. All real systems (processes) input energy and material and do work that produces a product and wastes (material and heat).

The decision processor entity in Figure 3 did not just jump into existence. It evolved from a precursor entity that actually received input from the work process entity and, at one time, presumably used that product for some other work process of its own. Recall that evolution depends on multiple copies of entities to generate variability without harming the basic entity plan. A really interesting trick in evolution has turned out to be that some copies (e.g. copies of entity B in Figure 2) are essentially redundant and become subject to separate selection forces. Thus one copy of entity B could diverge from the normal use of the product of entity A and begin to concentrate on feeding back signals to A relating to its “perception” of A’s product (remember B found A’s product useful in the first place to form the association). In other contexts this is the basis of speciation. Here it is the basis of differentiation (that will later show up in embryonic development).

The development of a specialized feedback mechanism that acts to regulate the work done in the work process entity (via an appropriate actuator) is the beginning of the cybernetic system, a system that uses feedback (at this level) to self-regulate the quality (quantity) of production so as to continue to fulfil its function. That is to supply a product that can be used by other type B entities (and others); that becomes its “purpose.”

HIERARCHICAL CYBERNETIC SYSTEMS

There is very little in the literature regarding the nature of a hierarchical cybernetic system (HCS). A Google search for the term “**hierarchical cybernetics**” (in quotes) turned up a mere 152 hits with the top ones being from my own work. A search for

A Systems Science Framework for Governance

“hierarchical **control**” turned up 246,000 hits. A quick survey of some of the top hits (some of which are also mine!) reveal a strong sense in which these are targeted for machine engineering; the word control being used in the context of machine control. For example a number of available papers address mobile robot controls where low-level “tasks” include sensory and motor reactive behaviours and high-level tasks include integration or blending of low-level behaviours and mission-achievement planning (Arkin, 1998, c.f. chapter 6). Other applications are for complex physical plants such as refineries and steel production. All of these types of applications share the property of relatively high levels of determinism in the lowest levels where classical control theory can be used with mathematical methods well understood.

The word “control” carries some baggage when thinking about social issues. For example a number of management theory researchers have long pointed out that people do not like to be “controlled” in the top-down command-and-control sense. Since I am interested in the governance of human organization/societies, I have largely abandoned the word control in favour of the more generalized *cybernetics* as it was conceived by Wiener (1950) and explicated as applying to living systems as much as to machines.

Lately interest in “distributed control” has also grown out of a concern for the top-down aspects associated with classically conceived hierarchical control. In engineering control systems for processing plants this is really another version of the hierarchical control model. In management the idea of distributed control is embodied in, for example, the “flat organization” favoured by some high-tech enterprises where they promote the notion of “local control” or decision making.

A major difference between machines and organizations/societies of humans is the degree to which non-determinism and non-stationarity play a role in the latter. The basic construct of cybernetics holds forth in all CAESs, that is, information flow integrates and acts to regulate the subsystems for the benefit of the whole. But there are differences in the kinds of decisions and the timing of such decisions that pertain to levels in a hierarchical system. Except in management science (and correspondingly in military organization science) the recognition of these decision types has not been adequately addressed, even in the engineering fields.

Namely, CAESs are hierarchically organized such that basic production operations form the lowest level. This level is populated by many different operational units, each attempting to produce a predetermined product (or service) of an essential quality that collectively lead to the final production of products (or services) for export to other entities. Each of these operational units uses the kind of feedback mechanism shown in Figure 3 to self-regulate or produce their special product in accordance with an “ordained” standard, the latter having been established as necessary to meet the final goals of the system.

But at some level of complexity (as defined previously) many sub-goals of many operations may compete for resources, or fall out of synchrony with others that depend on their products as inputs. The matrix of operational units requires a more global perspective and coordination in order to keep the whole functioning according to

A Systems Science Framework for Governance

requirements. An HCS “architecture,” then, consists of a base level of operations that are locally controlled by direct feedback mechanisms as in Figure 3. Above that level is a level of coordination control divided into two types of coordination. One form of coordination involves that between operating units. This is the logistic control orchestrated by logistical decisions. The other form involves coordinating the behaviour of the system (entity) with the dynamics of the external environment, especially those other factors and entities with which the entity interacts routinely, e.g. sources of food.

Natural Hierarchical Cybernetic Systems in Organisms

Living organisms offer some interesting insights into the hierarchy of cybernetic subsystems. They are not deterministic as a machine, though their level of stochasticity is not so high that too much error accumulates over the life time of the individual. If randomness dominated the processes then life itself would not be possible. But they are not subject to intrinsic non-stationarity which makes them a good starting point for analysis of natural HCSs. Living systems, from cells to organisms, use the hierarchical cybernetic system to mediate stochasticity and to take corrective actions in response to non-stationary stochasticity in their environments. In this section I review some aspects of the HCS model as it is recognized in living organisms.

Homeostasis

The capacity of critical mechanisms to maintain their operational parameters in living organisms is homeostasis. Since a tremendous amount of research has been done on this mechanism I will forgo descriptions in this paper. I note only that this is the basic mechanism of operational-level control that is ubiquitous throughout the living world and all scales (within cells, cell activities, tissues, and whole organisms). The notion of homeostasis in social settings has been explored as well (c.f. Richardson, 1991, p48).

Autopoiesis

Maturana and Varela (1973, 1987) defined life in operational terms as the ability for the system to self-create and maintain. A close examination of several living systems, including the metabolism within cells and physiology of whole multicellular organisms reveals that a number of mechanisms have been evolved that supplement the basic homeostatic cores of productive work with networks of coordination functions that receive information from the basic cores and respond to demands for resources (including repair and replacement of “machinery”) but in a way that balances the needs of many separate cores. For example the expression of genetic information involves not only the transcription of DNA into messenger RNA molecules, but a demand-driven cadre of low-weight RNA molecules that can interfere with translation by blocking the mRNAs before they can attach to the ribosomes. Similarly other epigenetic mechanisms have been found that enhance or promote specific transcriptions in response to internal needs of the cell.

The general autopoietic matrix that envelops the basic production activities of metabolism acts as a kind of logistic coordination framework to ensure continued smooth and balanced functioning.

Adaptive Mechanisms

Cells have evolved a large array of membrane channel proteins that can be activated by the presence of low-weight molecules in the medium outside the cell. The opening or closing of said channels activate a variety of internal signalling mechanisms that lead ultimately to changes in the outward behaviour of the cells as well as some internal changes that may persist. For example neurons, in particular their synaptic junctions react to neurotransmitter molecules that may cause the synapse to generate a membrane discharge that could lead to the neuron generating an action potential, thus signalling other neurons. Moreover, if such discharges happen often enough, especially in conjunction with other incoming signals, the sensitivity of the synapse may be changed and it responds adaptively to future incoming signals. The neuron as a whole will behave differently with different overall consequences in the future as a result. This is a micro-scale version of coordination with external events that, in the case of the whole animal, results in macro-scale coordination (i.e. learning to react to whatever caused the change).

Evolvable Mechanisms

The replication of the genome of cells, through mitosis, or the whole organism, through meiosis, is fairly well understood as the copying mechanism that provides an opportunity for variation, through mutation or cross-over. What is somewhat less understood, but subject to intense investigation now, is that the variation generation rate may actually be under some kind of control mechanism that promotes the mutation rate in sections of the genome that affect traits that may create more fit variations. In the case of gametogenesis the effect is to generate variance in subsequent generations of individuals as fodder for selection, but in general body mitosis it can lead to stereotypical cancers. Various other epigenetic mechanisms have been discovered that have heritable properties suggesting that some species can acquire properties (not new traits per se) that are passed on to offspring.

Evolvability appears to increase the chances that at least some individuals in a population will develop capabilities that make them more fit in light of long-term environmental stresses. It might be considered as a form of intentional evolution with strategic consequences for the genus.

Architecture of a Hierarchical Cybernetic System

Overview

The general HCS architecture for all CAESs is shown in Figure 4 below. It is a three layered architecture based on types of decisions involved in the overall governance of the system. The division in layers is also based on time domains over which decisions and actions need to be made. Within layers there may be a range of decision scopes and time scales. For example in extremely complex system with many different sub-types of operational cores there can be a sub-hierarchy of coordination modules, e.g. a supra-coordinator of coordination modules such as is found in large organizations with large departments and sub-departments. Each sub-department has a coordinator (middle

A Systems Science Framework for Governance

manager) that reports to a higher-level coordinator, e.g. the accounting manager reports to the controller.

Time Scales and Decision Types

Operations level decisions and actions are generally done in what we can generically call “real-time.” The actual range of such decisions and actions is variable but entails decision agents operating in the domain using the same time constants as characterize the work operations being done. The decision types are corrective. That is they are the basic homeostatic decisions based on keeping outputs in conformance with standards (ideals). This is the type of control decisions normally discussed in treatments of cybernetics.

Logistical decisions are taken over longer time scales in general and rely on aggregated (and often weighted) data coming from the operations level. Time averaged performance data is used with more complex decision models, such as optimization of overall performance requirements under constraints with disturbances accounted for.

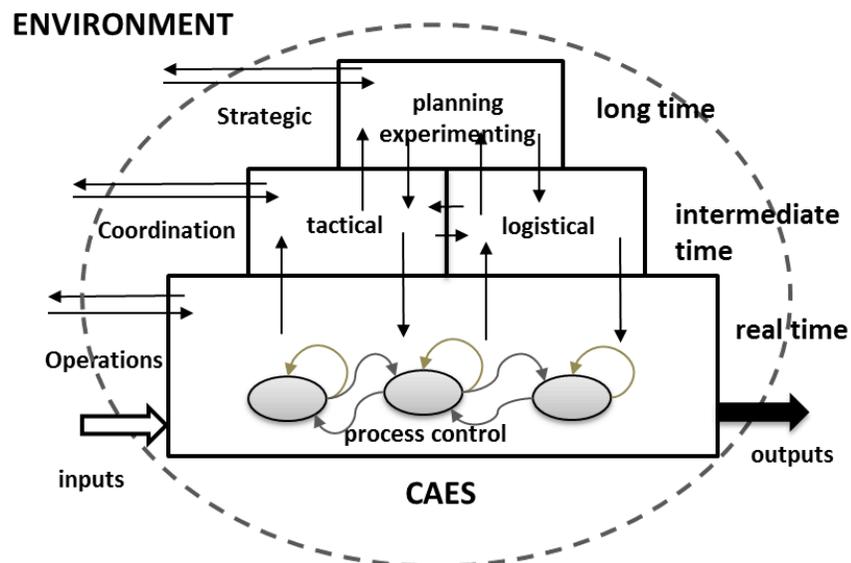


Figure 4. The HCS architecture is a three layer structure divided by decision types and time domains.

Tactical decisions operate across a range of time scales generally longer than operations, though it should be noted that many tactical decisions have to be carried out by operations level processes. For example, a brain may decide to go from the animal’s current position to another location to see if food is available. The planning of actions is done in the tactical part of the brain but the execution is an operational level action taking (muscles and senses).

As I will argue below, the capacity for real strategic thinking seems to be confined to humans (and possibly pre-human precursors) and human organizations. For all of the rest

A Systems Science Framework for Governance

of biology goals are selected for species by evolutionary processes. Most animals have limited autonomy in this sense. But when strategic capabilities are part of an individual entity's decision processing the types of decisions made involve more comprehensive evaluations of environmental conditions and their evolution into the future. New goals can be generated and as long as they conform to requirements of the biological mandates they lead to new behaviours. The time domain for strategic goal setting is generally long with respect to the "average" life of the system but varies with types of systems.

Communications Channels

Communications between subsystems involves the emergence and selection for low-energy "channels" and information encoding schemes. Channels can be highly diverse and carry very different kinds of messages depending on which level in the hierarchy. Additionally channels between levels are needed. The senders and receivers of messages share a common protocol that determines what encoding and decoding of messages are allowed (and are meaningful).

The key issue regarding communications between subsystems within and between levels is that messages travel rapidly compared with material and energy flows and work processes. Myelinated nerve fibres or electronic wires convey messages far more rapidly than the time scales over which the processes they affect operate. This is even true for chemical messengers travelling through the blood stream (like adrenaline).

Computations and Decision Agents

The architecture makes provisions for the kinds of decisions and the supporting computations that need to occur at the various levels. Generally speaking the computational loads increase at higher levels. Logistical computations require more elaborate control models (e.g. optimization models compared with simple reaction models, like a thermostat) and more memory space. In evolutionary terms, acquiring coordination control is expensive so the benefits must outweigh the costs. Presumably there is a level of complexity that when reached requires specialized coordination agents with their extra costs in order to obtain a gain that exceeds those costs. More details on decision agents is given below.

Operations Level

Operations level management is largely handled by local feedback control as in Figure 3 above supplemented by cooperative information flows as in Figure 2. In more elaborate versions the cooperative information flows form what we would call markets. In (Mobus, 2016, in early development) I provide a full accounting for how market mechanisms act at this level. I also cover the dysfunctions that obtain with increased size and complexities of markets when coordination supplements are not developed (or when markets are not regulated!)

Coordination Level

Tactical and logistical decisions are both examples of coordination among diverse and disparate processes and other entities. The multitude of internal subsystems and processes need to be coordinated so that they succeed in collectively producing the final objectives (products) gives rise to a nearly exponential explosion in communications and computation load.

Figure 5 shows a simple representation of the emergence of the coordination level agents that will assist the lower operations level processors keep in synchrony and balanced insofar as their production rates and qualities are concerned. This diagram also shows the relations of operational level controls and the coordination level interfaces.

Figure 6 shows the communications complexity explosion as the overall system grows in size and complexity. At some point the communications and computation loads on the first layer of coordination agents becomes such that a new superior layer emerges to coordinate the coordinators.

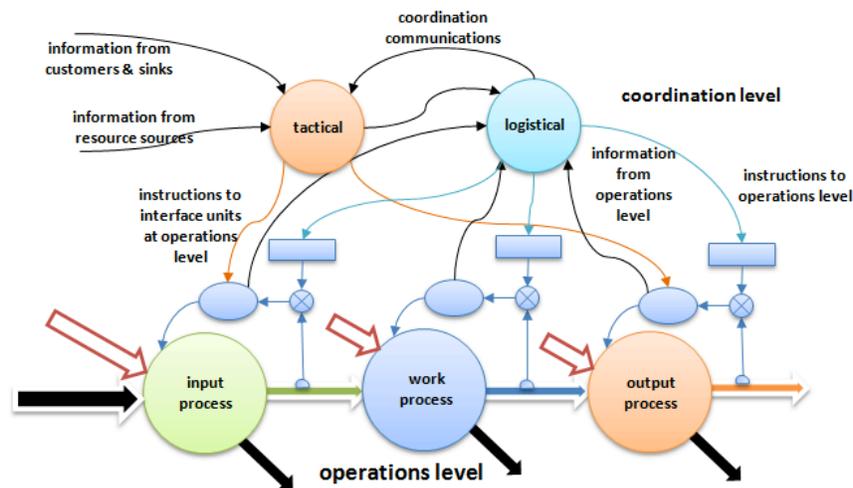


Figure 5. The coordination level agents (tactical and logistical) act to coordinate the behaviours of the operations level work processes. Here specific input and output work processes are coordinated by the tactical agent but are also assisted by the logistical agent insofar as interfacing with the other internal work agents is concerned. In a realistic model the central blue work process would actually represent many sub-processes and the logistical problem would be substantial

A Systems Science Framework for Governance

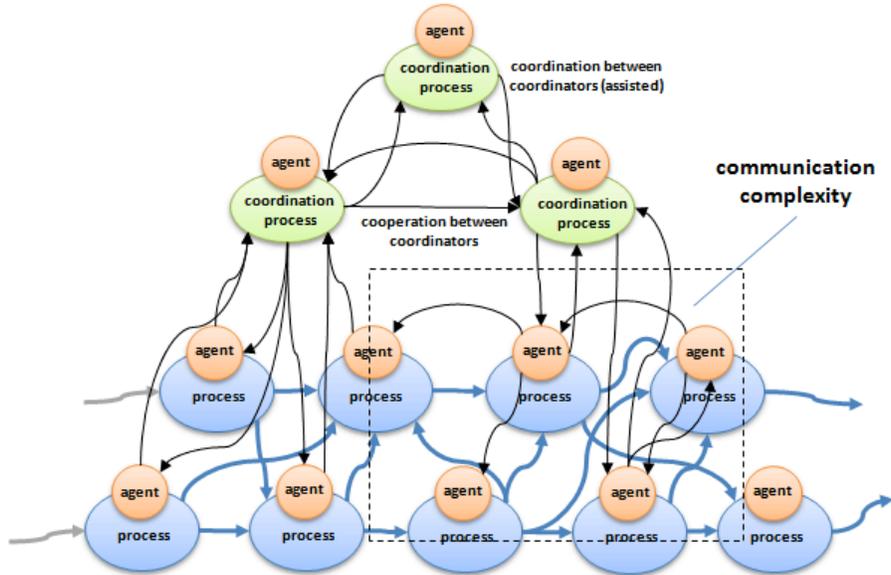


Figure 6. When the whole system becomes large and complex in terms of the number of work processes that need coordination the computational and communications load (if every communications channel within the dotted lined box were implemented) become too much for affordable coordination agents.

In Figure 6 the problem of computational and communications load is solved by the emergence of a higher-level (or a new layer) of coordination agent(s), say an agent that helps coordinate the lower coordination agents. Of course this is precisely what we see in large organizations when middle management deepens in response to increasing complexity (e.g. size of departments). But it is also the case in living systems. For example, in cellular metabolism there are a raft of “middle managers” acting to coordinate the DNA transcription, RNA transmission, polypeptide translation, and protein assembly processes. And there are higher middle managers that help coordinate those, e.g. the iRNA molecules sent to “interfere” with some aspects of the process under specific conditions. Biologists are discovering more of these systems coordinating subsystems quite frequently it seems.

Strategic Level

The Commander-in-Chief of the military, the CEO of a corporation, and the president of a university are not only the top slots in an organization chart, they are the people tasked with making decisions about what to do in the future to best position the country or organization relative to opportunities and threats that will have impact on the system in that future. Their job is to think about the world in which their system operates and will do so in the future. They need to set in motion plans for strengthening weak areas of the system and acquiring capabilities that will be needed to meet the future demands. As argued above, it seems that only human organizations and states, as well as human beings

A Systems Science Framework for Governance

themselves, have evolved this level of agency. As I will consider below, the competency of this level is still in serious question.

All other biological entities have evolved ways of living, phenotypes and behaviours that constitute their “strategies” in moving into the future. They do not “construct” such strategies intentionally. Only human agents show any capabilities in that domain. But as we will see, that capability is still very weak and is not always operative for individuals or organizations. It is very weak in terms of social governance.

DECISION AGENTS

Generic Decision Agents

Every kind of cybernetic process requires that there be a selection of a control signal value out of a range (or set) of values that will be sent to an actuator based on a mapping from the state of the system and the current conditions of the environment to that signal (value and valence). The mapping could be as simple as a proportional function and mechanically decided (e.g., the thermostat) or it can be a multivariate, non-linear function that requires considerable computation. The mapping might involve probabilistic characteristics. Tactical and logistical decisions tend toward the latter realm. Whatever the mapping is, it is derived from a control model of the physical process.

Until the age of the computer and what we now call cyber-physical systems the more complex sorts of coordination decisions were left to human beings who could operate with elements of uncertainty and non-linearity and provide reasonable estimations of what needed to be done. At present airplanes still need pilots but our theories of control continue to improve.

Figure 7 shows the basic schema of a decision agent.

The decision agent is essentially a computational engine that uses report messages from lower level units regarding the state of the subsystems over which it responsible, a decision model that captures the form and dynamics of the decision type (e.g. use a linear programming model to compute optimal outputs), and sends various kinds of command signals to the lower units to spur actions. One example of this for a logistics agent is tuning operational level set point values (Figure 3, the “ideal” value) to adjust flows through the work processes as explicitly shown in Figure 5. If the model being used can be modified on the fly, or upgraded as a result of learning, then the agent is adaptive.

A Systems Science Framework for Governance

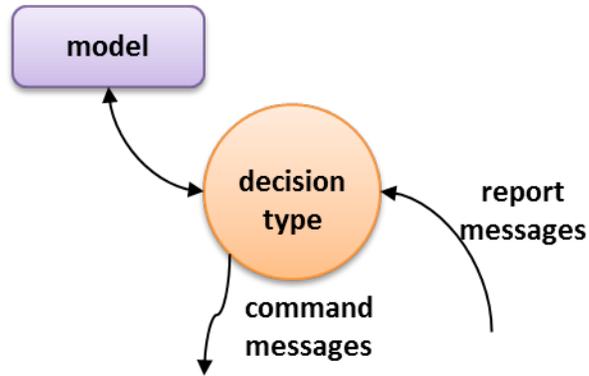


Figure 5. A general schema for a decision agent based on type of decision. In all cases the agent receives messages from units lower in the hierarchy (or sensors) and uses a decision model for the type to generate command messages that are sent back down the hierarchy.

Decision models can be rule-based (algorithmic) or heuristic. They can be learned or designed depending on complexity, stochasticity, and non-stationarity. In the latter case models are used to approximate appropriate command signals. In the case of tactical models they are used to anticipate the behaviours of external entities and environmental states in order to reduce the costs of merely reacting to changes (Mobus, 1999).

In the context of social governance the agents range from individuals sitting in particular positions of “authority” and “power.” A fundamental problem with human decision agents is the abuse of these concepts that I will discuss briefly below. But human decision agents can also be groups, like committees, panels, etc. We should grapple with the nature of human decision agents because they add a level of complexity to the problem of governance of social systems that is not often appreciated. Humans may be strongly motivated to cooperate, but they are also autonomous CAESs on their own terms.

Human Decision Agents

Adam Smith did not just leave us with the idea that free markets would suffice to provide for the good of the whole society (Smith, 1776) in his most referred to work. In a prior work, *The Theory of Moral Sentiments* (Smith, 1759) he painted a picture of Man as naturally being concerned with his neighbours (and those with whom he did business) and though selfish in some regards, nevertheless could not live as a purely selfish agent. I have always viewed Adam Smith as a keen observer of human nature and behaviour, so I am in agreement with this idea.

In other work in which I am engaged I have been investigating the psychological construct of wisdom and its processing basis in the brain (Mobus, 2015, in preparation). I call the latter sapience. It is similar in psychological terms to intelligence in that some brains seem to possess greater capacity than others. The key aspect of wisdom (and the sapience that supports it) is veridical higher-order judgment regarding complex

A Systems Science Framework for Governance

(“wicked”) social problems that addresses solutions that serve the best interests of the greatest number of people involved.

A Decision Agent with its Own HCS!

The situation that leads to problems for human social governance is the simple fact that each individual human is an HCS on its own accord. Humans are autonomous agents and are often motivated by selfish factors. The problem, as I will briefly discuss below, is that humans have a strategic thinking ability but have serious limitations on its capacity. Thus all organizations/societies in which humans are the primary decision agents will be subject to irregularities that have their origins in individuality. The degree to which human agents can process strategic decisions will have an impact on group dynamics and successes. If an individual is highly strategic then they will be thinking about the good for the group and not just themselves. This is because a true strategic viewpoint for an individual admits to the fact that what is good for the group is good for the individual (in general). Humans have a capacity to think strategically for their own benefit but it will be limited to more “immediate” gains for themselves. It will not necessarily encompass the long term or the wider scope of the group.

Strategic thinking is necessary but hardly sufficient. Above I said that a “true” strategic viewpoint will encompass the good of the group. What makes it “true”? One counter example of strategic thinking that is not really necessarily good for the group is when it is motivated by ideology rather than wisdom. Certain leaders have demonstrated time and again an ability to work out strategic moves for society believing fervently that their ideological views constituted what was best for the people only to have the outcomes make their people worse off than before. For example the neoliberal capitalist ideology that is sweeping through many otherwise democratic, even socialist-leaning states, is leaving a trail of what is known in network theory as a scale-free network structure for wealth accumulation or the phenomenon of “the rich get richer” (Barabási, 2002). The prevailing claim is that “...a rising tide lifts all boats,” is taken as self-evident as is the claim that “...the rich are the jobs creators.” These are not empirical results as much as faith-based ideas. Yet they have caught on in governments around the world and are justifications for the highly biased wealth distribution we are seeing in practice. The point is that some very clever people have developed strategic plans for how to move that agenda forward. But those people cannot be called wise or seriously having the good of all in mind. It is not established that the wealthy actually create jobs, for example.

Wisdom requires a few more components in order to be realized in a human individual (Mobus, 2015, in preparation). Among them is a clear moral sentiment that motivates concern for all (for example actually seeing to it that all boats are rising!) Another is comprehensive systems thinking that allows one to contemplate the long-term consequences of various actions (e.g. policies) that might be undertaken.

Mobus (2015, in preparation) provides evidence that the majority of human beings have relatively weak capabilities in these various components. They are present, but not well developed, or more likely, sufficiently evolved to assure the adequate development of wisdom with age.

A Systems Science Framework for Governance

Figure 8, below, depicts a human decision agent (adaptive and evolvable) showing the HCS components.

Bounded Rationality and Bounded Wisdom

The brain capacity for sapience is essentially subject to the same kinds of constraints as for intelligence. That is, as Simon (1998) recognized about intelligent decision-making, so also wisdom is bounded by processing capacity and time. Intelligence, creativity, and wisdom are constrained by brain processing capacities, and those are strongly influenced by genetics. Wisdom is statistically rare (Sternberg, 1990a). And it only emerges as a function of age and experience. It does not directly affect rational decisions, but influences decision processing through intuition and subconscious processing (Mobus, 2015, in preparation).

Human decision making is a strange mix of rational, emotional, and intuitional driven factors. The rational seems fairly rare if the results from psychology stand up (Kahneman, 2011). Emotionally motivated (and directed) decision making seems a good deal more prevalent in human affairs. And that leads to a fundamental conundrum. As components in an HCS, humans are unreliable. They make mistakes in judgment that have consequences when they are in positions as decision agents in a governance system. One cannot help but wonder if the magnitude of the consequences, especially the negative ones, scales with the size and complexity of the system being governed.

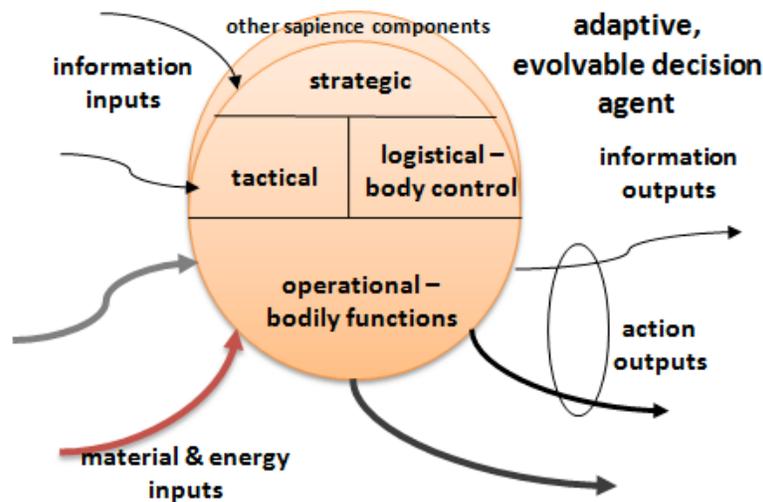


Figure 8. A human decision agent is an HCS and thus an autonomous agent. The efficacy of decisions that are group-productive depends on the strength of the strategic level of decisions supported by other components of sapience.

HUMAN SOCIETY IN THE ECOS

Any contemplation of a systems-based governance of human society must start with the recognition of that society as a subsystem in the Ecos. *Homo sapiens* evolved in the Earth ecosystem and the species' various societies have always been embedded in that system. Figure 9 illustrates a very macroscopic view of society in the Ecos.

The Ecos is an effectively closed system with respect to material flows, at least as far as human time scales. The energy flux it receives is essentially in a steady-state (fluctuating within fairly narrow ranges) and the availability of free energy has, for most of the history of the planet, depended upon the efficacy of photosynthesis, the evaporation of water and the driving of climate and wind. Ancient photosynthesis was responsible for sequestering a fair amount of energy in the form of organic matter that would turn into fossil fuels. Those fuels are now being burned by human society to accomplish exosomatic physiological support work, i.e. our technological cultures and releasing waste heat and combustion products faster than the Ecos' natural rates can accommodate.

If human society is a CAES then one might reasonably look for regulating feedback loops that allow the system to evolve in balance with available resources without destroying the other subsystems that, for example, provide various recycling services to the Ecos.

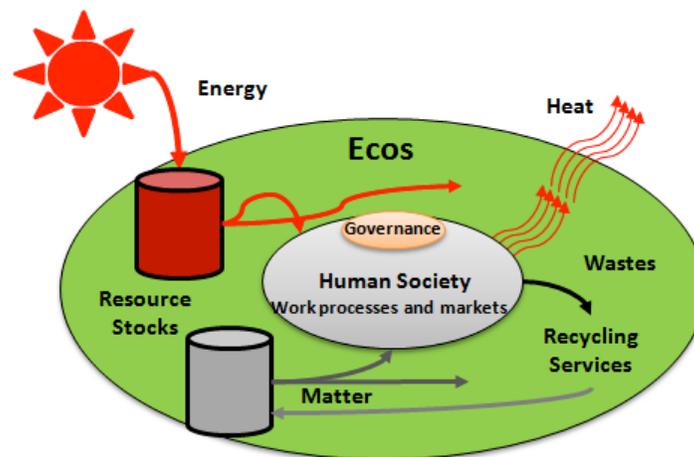


Figure 9. Human society is just one of many subsystems of life, geology, etc. that are contained fully within the Ecos. Human society, at present, occupies a larger than reasonable space within the Ecos and generates more entropy than the ordinary degradation by other Ecos subsystems of energy flows from the Sun.

Human Societies as CAESs – A Program of Investigation

We can view society as a hierarchical system of self-similar (fractal) units at increasing levels of organization and complexity. If we start with the individual human as a kind of “atomic” CAES entity we can characterize many levels of organization from there up to society as a whole (individuals as members of a corporation, that as a participant in the

A Systems Science Framework for Governance

marketplace, that as a subsystem of a region, etc.). But society is also embedded within the entire Earth system – the Ecos. The latter is effectively a materially closed system but is open to energy flows and when formed originally was relatively speaking “simple.” Thus over the life of the Ecos to the present it has had an abundance of simple components that gave rise to life and that life evolved to a high level of organization with the advent of humans and their societies. The Ecos is a CAES by virtue of the fact that it started with very high potential complexity and has evolved to the kind of realized complexity we see today. So long as free energy is available to do work and generate increasing organization the system will continue to evolve. But we can imagine a time when the flow of total energy available to the human society either comes to steady state again or declines such that free energy per capita is greatly reduced.

The Natural Human Social Governance System

A Starting Point

The social nature of hominids (and indeed most primates) is well established. The evolution of increasingly complex social interactions is also reasonably well understood. Humans represent the epitome of social evolution and by the early Holocene human tribes had formed with natural hierarchical structures based on the kinds of decisions (strategic, coordination, operations), the relative wisdom of the decision makers, and the specialization in skills and decision making talents. Figure 10 shows a distilled (and simplified) schematic of a tribal organization along these dimensions.

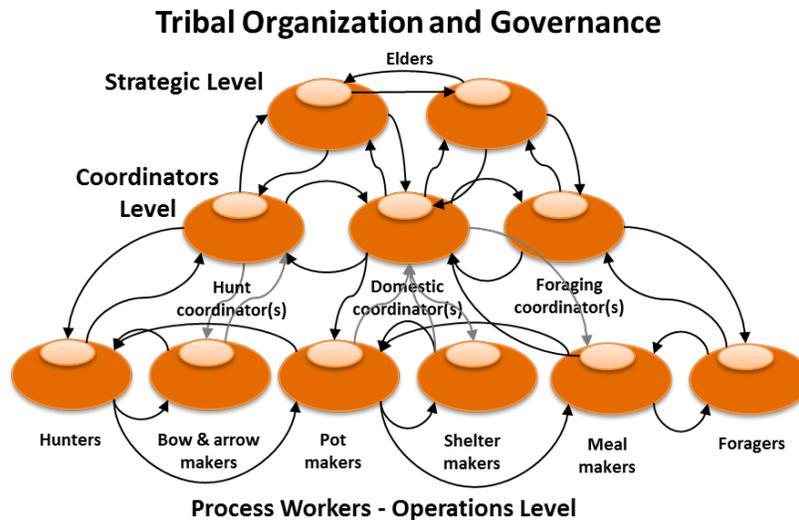


Figure 10. Humans evolved a social structure that reflected the HCS architecture during the Holocene era. The tribal organization served human evolution through group selection. The roles labelled give an idea of how the work and decision processing were covered. Individuals might very well have filled multiple roles in any one tribe.

A Systems Science Framework for Governance

Given that this kind of organization is also the form of governance that humans first evolved we can use it as a starting point, or reference model, in our examinations of governance throughout history and in its modern forms.

Decision Models for Social Governance

Humans have experimented around with a number of governance macro-scale models based on a large number of factors that were as much the product of speculative thinking as anything. While I assert that a governance system is fundamentally structured as an evolvable HCS these many other factors help shape the actual form of communications and all of the lower-level decision models that are developed at all levels.

Starting with basic philosophical considerations, e.g. human nature is Hobbesian or Rousseau(ian), and concepts of the sources and uses of physical wealth (Marxism or capitalism or some other ‘ism’) various ideologies of governance have overlain the fundamental HCS framework with specific ‘mechanisms’ for establishing communications channels (e.g. markets) and decision models (e.g. dictatorships vs. democracies). Constitutions are the grand overarching models for nations. Bylaws play a similar role in smaller organizations. Civil laws, contracts, ethical codes, market rules, and various hand-shake agreements are all ways of establishing cooperation and coordination mechanisms in the social order. Moral codes, such as the Golden Rule, originate in the moral sentiments from the limbic system in the brain.

Many different forms have been invented and tried for governing every scale of organization. The most impactful have been the forms applied to national and regional geographies. Civilizations have grown, empires established and eventually fallen, and every conceivable kind of coordination transaction has been made, from strong and inhumane coercion, to gentle prodding, to outright rational conversation (or debate). To date none has produced a stable, sustainable society. Not even the current neoliberal capitalist (democracy or not) has given us a social system free from unfairness and living up to the claims of what has been billed as the “American Dream.” Why not?

Conclusion and a New Approach

We humans have tried many approaches to governing our social systems. We’ve put a lot of faith in ideologically-derived mechanisms (e.g. free markets solve all problems) and some of these have worked a bit here and there, while also failing in other aspects. The free markets have not been able to solve most of the equity issues our present society presents to us. So far it has been a process of learning a little from the past, at least recognizing the parts we didn’t like so much, and then trying to produce something new that is a better form of governance (where better means reaching some of the objectives of the freedom and pursuit of happiness). Some would argue that things have gotten better and fairer, that the American Dream and American ideals (e.g. torture is wrong) are far better than feudalism. Maybe. The fact that we are talking about an Anthropocene suggests that maybe humans have not yet established a governance form that keeps us in balance with the Ecos while bringing about an equitable and reasonable standard of living.

A Systems Science Framework for Governance

The Ecos is not getting bigger. It is essentially fixed in terms of material resources. The flux of high-grade energy is also essentially fixed; energy is ultimately flow limited. These are the constraints that our society must live with for the foreseeable future. Just as the internal governance of the body brings physical growth to a halt at some point, so too the governance of society must establish a steady-state, sustainable size and consumption rate. The society of humans must become an entity just as the many transitions in biology before us have done. A systems science approach to the intentional design of governance for that purpose offers one possible route to that goal.

REFERENCES

- Albus, J. S. (1996). "The Engineering of Mind". *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*. The MIT Press, Cambridge MA. Available on-line from Google Books at: <https://books.google.com/books?hl=en&lr=&ie=UTF-8&id=V3pksEEKxUkC&oi=fnd&pg=PA23&dq=The+engineering+of+mind&ots=sFjNGvuR8m&sig=Hrvf8uN732vQVBhgJjrohsUtfqI#v=onepage&q=The%20engineering%20of%20mind&f=false> [page 26 – Complexity management through hierarchical architecture]
- Alkon, D. L., (1987). *Memory Traces in the Brain*, Cambridge University Press, Cambridge.
- Arkin, R. C. (1998). *Behavior-Based Robotics*, The MIT Press, Cambridge MA.
- Barabási A-L. (2002). *Linked: How Everything is Connected to Everything Else and What It Means for Business, Science, and Everyday Life*, Penguin Group, New York.
- Beer, S. (1980). *The Heart of the Enterprise*, Wiley, New York.
- Beer, S. (1981). *Brain of the Firm*, Wiley, New York.
- Bourke, A. F. G. (2011). *Principles of Social Evolution*, Oxford University Press, Oxford UK.
- Corning, P. (2010). *Holistic Darwinism: Synergy, Cybernetics, and the Bioeconomics of Evolution*, University of Chicago Press, Chicago IL.
- Calcott, B. and Sterelny, K. (eds. 2011). *The Major Transitions in Evolution Revisited*, The MIT Press, Cambridge MA.
- Kahneman, D. (2011). *Thinking Fast and Slow*, Farrar, Straus, and Giroux, New York.
- Lovelock, J. (2006). *The Revenge of Gaia: Earth's Climate Crisis & the Fate of Humanity*, Basic Books, New York.
- Maturana, H. and Varela, F. ([1st edition 1973] 1980). "Autopoiesis and Cognition: the Realization of the Living", Robert S. Cohen and Marx W. Wartofsky (Eds.), *Boston Studies in the Philosophy of Science* 42. Dordrecht: D. Reidel Publishing Co.

A Systems Science Framework for Governance

- Maturana, H. R. and Varela, F. (1987). *The tree of knowledge: The biological roots of human understanding*, Shambhala Publications, Boston, MA.
- Meadows, D. (2008). *Thinking in Systems: A Primer*, (D. Wright, ed.), Chelsea Green Publishing, White River Junction, VT.
- Mobus, G. and Kalton, M. (2014). *Principles of Systems Science*, Springer, New York.
- Mobus, G. (2015, in development). *Sapience: Understanding the Human Mind from a Systems Science Perspective*, Draft available to interested reviewers; send e-mail request.
- Mobus, G. (2016, in early development). *The Systems Science Approach to Governance for Sustainable Society*.
- Mobus, G. (1999), Foraging Search: Prototypical Intelligence, *The Third International Conference on Computing Anticipatory Systems*, Liege, Belgium, August, 1999. Available on-line: <http://faculty.washington.edu/gmobus/ForagingSearch/Foraging.html>
- Morowitz, H. J., (2004). *The Emergence of Everything: How the World Became Complex*, Oxford University Press, Oxford UK.
- Richardson, G. P. (1991). *Feedback Thought in Social Science and Systems Theory*, University of Pennsylvania Press, Philadelphia PA.
- Simon, H. A. (1998). *The Sciences of the Artificial*, Third Edition, The MIT Press, Cambridge MA.
- Smith, A. (1759). in Knud Haakonssen, ed. (2002). *The Theory of Moral Sentiments*, Cambridge University Press.
- Smith, A. (1777) [1776]. *An Inquiry into the Nature and Causes of the Wealth of Nations*. University of Chicago Press.
- Smith, J. M. & Szathmáry, E. (1995). *The Major Transitions in Evolution*, Oxford University Press, Oxford UK.
- Sober, E. and Wilson, D. S. (1998). *Unto Others: The Evolution and Psychology of Unselfish Behavior*, Harvard University Press, Cambridge MA.
- Sternberg, R. J. (ed.) (1990a). *Wisdom: Its Nature, Origins, and Development*, Cambridge University Press, New York.
- Sternberg, R. J., (1990b). "Wisdom and its relations to intelligence and creativity", in Sternberg, R. J. (ed.) (1990a, pp 142 - 159).
- Tomasello, M. (2014). *A Natural History of Human Thinking*, Harvard University Press, Cambridge MA.
- von Bertalanffy, L. (1968). *General Systems Theory: Foundations, Development, Applications*, George Braziller, New York.

A Systems Science Framework for Governance

von Neumann, J. and Morgenstern, O. (1944), *Theory of games and economic behavior*, Princeton University Press, Princeton NJ.

Wiener, N. (1950). *The Human Use of Human Beings: Cybernetics and Society*, Avon Books, New York.

Wilson, E. O. (2013). *The Social Conquest of Earth*, Liveright, New York.