Chapter 10 – The Agent Model

Abstract

Purpose of this Chapter

Agents take[1] decisions and take action. In this appendix we will describe the model archetype for a generalized decision agent. This model will be revisited in the next appendix covering the model archetype of governance. Decision agents are the actors in governance and management of CAS/CAESs. Their agency is in the fact that they can sense/measure relevant parameters, compute via a decision model what action should be taken, and then send signals to the appropriate actuators to effect the decision. After describing the basic architecture of the agent model we will provide several examples from biology, cyber-physical systems, and social systems.

Agency and Agents

In our investigations of CAESs we inevitably find ourselves needing to understand agency (the power to affect other components in the environment) and the nature of agents as the entities that take decisions and take actions; in other words, have agency. Within even the most complex governance architectures agents at different levels of 'authority' also have different degrees of 'autonomy,' or the freedom to make choices or take decisions[2]. This correlates with the complexity of the decision domain in which they operate.

The nature of an agent, its degree of autonomy, and its agency will be examined in more detail since these are the most critical components of a governance system.

The question of who benefits from the actions of an agent is a key element. Does the actions of an agent benefit its self, benefit other entities, or benefit the whole system through the network of agents in a governance architecture? We argue that, in fact, the answer is all three will benefit

[1] We will be using the term 'take' in the way Europeans refer to decision making but generally interchangeable with the term 'make'. At least one author differentiates between 'taking' and 'making' of decisions where the former is an instantaneous event – selecting from a set of decision pathways – whereas the latter is a process that may include constructing the pathways from analysis. See the blog article: https://www.excitant.co.uk/decision-making-and-decision-taking-they-are-different/ for a discussion. Accessed 5/30/2018.

[2] See the Wikipedia article: https://en.wikipedia.org/wiki/Autonomous_agency_theory for background. Accesses 01/01/2017.

The Basic Agent System

In this section, we will consider the systems model of an agent and situate that agent in a larger CAES in which its decisions and actions affect the operations and sustainability of the whole larger system. It is agents that constitute the workings of governance and management processes so understanding what they are, how they work, and particularly what happens when they make mistakes is essential.

Figure 10.1 shows a basic and universal model of an agent. The operational elements of an agent are: a computational engine[3], a decision model, and an experiential memory. The latter may have been built from learning in higher animals, organizations, and artificial agents (e.g. robots) or the "memory" may be built into the structure of the system through evolutionary modifications to brain circuits that control behavior, i.e. instincts.

In cyber-physical systems (e.g. robots) the three roles are easily seen to be separated, the decision model (i.e. the program) and experience memory (variables and data bases) are both contained in the computer memory; the computational engine is, of course, the hardware. In brains, however, all three components are highly integrated in the nature of neural computation involving networks of interconnections between neurons; the processing, encoding of memories, and the models are all embodied within the networks of potentiated connections.

In the human brain, functions of these components at higher levels of organization are distributed among multiple specialized modules in the various cortices (see the section below on the human brain as an HCGS). Each module uses the same agent model, however, but working on a very specific set of inputs (from other modules) and producing very specific sets of outputs (going to other modules). The brain has been characterized as a network of agents. Minsky (1986) described the mind as a society of simple agents that acted almost like a democracy (each agent voting, so to speak, for a desired solution to a problem.

The agent is given agency, the power to affect the necessary response to the environmental conditions, when its output signals command an actuator having the necessary power to affect changes on those conditions (requisite variety, Ashby, 1958).

Agents are supplied motivations and goals from sources outside the basic agent model. For example, hunger is a drive arising in the hypothalamus that will drive an animal to seek and eat food[4]. The brain centers responsible for tactical control for food seeking (e.g. foraging) are activated and the animal goes into seeking mode. Human agents often receive motivations from

[3] As used here and explained in Mobus & Kalton (2015, Chapter 8) computation broadly covers any process that transforms input messages into output messages. This includes the traditional notion of the Universal Turing Machine running algorithms but also biological signal processing and activation working in the domain of chemical processes.

[4] See the Wikipedia article: https://en.wikipedia.org/wiki/Hunger_(motivational_state) for background. Accessed 7/15/2019.

other humans or, for example, from social norms. In our discussion below of human beings as agents we will revisit the role of motivation and goals as they pertain to decision making under conditions of high autonomy and intentionality as expressed in our species.

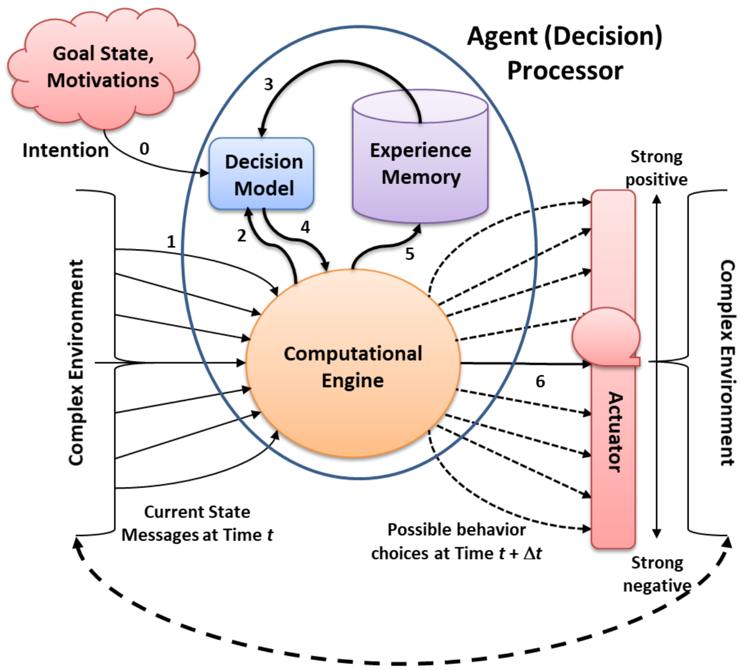

Fig. 10.1. A basic and general agent model includes a computational engine, a decision model, and an experience memory. Agents are supplied goals or motivations from external sources (represented here by a “cloud”). Sequence of events: 0) The decision model is provided with an intention (goal or motivation to attain a particular state of the world); 1) messages arrive from a complex environment providing state information at time *t*; 2) the computational engine uses the decision model to compute the decision; 3) the decision model uses prior experience from memory so that; 4) the computational engine can complete its work to; 5) update the experience memory (in the case of a learning agent) and; 6) make a choice of action out of many possible (dashed arrows). The output controls the actuators that affect the behavior that acts upon the complex environment at time $t + \Delta t$. Not shown is the energy flow inputs to the processes and the actuator.

What the figure does not show is the situation of a complex agent, which is an agent having many degrees of freedom and multiple variables to manipulate. In the next appendix we will tackle this issue more directly as a problem of governance of the agent as a system. In anticipation of that discussion we mention only that agents themselves may be sufficiently complex that they are organized as a hierarchy of agents that manage the whole agent/agency process. For example, consider the office of the CEO of a large corporation. The president is the chief strategic manager of the company but is also in charge of running a whole ‘department’ devoted to the various management functions needed to collect and process data and interpret the results in terms of strategic planning. The department is thus an agent for the whole company but it is composed of many sub-agents reporting to a 'master' agent[5].

Decision Making

The process of making a decision and activating action is based on the model in Figure C.1. The nature or type of decision determines what sort of decision model will be appropriate (see Appendix D for a description of decision types relevant to governance). The purpose of the decision at any time $t+\Delta t$ is dependent on the data collected regarding the state of the external 'complex' environment at time t along with relevant experiential memory[6]. The agent must interact through its motor outputs, collectively producing the behavior. The model that connects inputs with appropriate outputs, known in engineering cybernetics as the transfer function, and some form of context as contained in an experiential memory are used in the computation producing the behavior. Thus, the environment is altered by the action and gives rise to a new state to be measured by the input transducers (sensors) at a time later that includes delays for the computation and action (Δt) and latency for discernable changes in the state variables of the environment, a value that is dependent on many factors and in general will itself be variable. We assume constant time for the decision process itself in simple cases, but this is not the general case as complexity of the data and memory increase.

As we go up the hierarchy of complexity in concrete systems we need to increasingly take into account the role of stochastic process. This applies not only to complex environments but also to the complexities of the elements in the agent. We tend to think of digital computers as being completely deterministic, but even using digital computers in an agency of networked agents (called system of systems, SoS, in highly complex engineered systems) or distributed computation in cyber-physical systems, the vagaries of noisy communications, time lags, and data inconsistencies in distributed experiential memories (distributed databases) can introduce stochasticity into the general system. In living systems, chemical reactions are inherently stochastic even if constrained directionally by, for example, the specificity of enzymatic reactions.

The role of experiential memory varies depending on the degree of autonomy and the capacity for learning. The diagram in Figure 10.1 shows the various components of the decision process separated. In the case of, say, computer-based control systems this is generally the case. The decision model itself needs to include provisions for incorporating experience memory (e.g. state variables that are changed as a result of prior decisions) allowing its own structure to be

[5] The use of the term "master agent" is not to be taken literally. As will be shown in Appendix D (next) the idea of an agent higher up in the decision hierarchy is not that it need be a command-and-control master, but, rather, a coordinator of the activities of sub-agents.

[6] Those familiar with the concept of a finite state machine will see some resemblance. However, the general model of an agent does not require a finitude of states of the system (or of the inputs).

modulated by experience. In using computers for control systems this is a notoriously hard problem in machine learning. However, in the brains of mammals (and birds) the three functions are built into a single structure, the architecture of the neocortex (palladium in birds) in which adaptive memory encoding and computation are integrated. This is thought to be the nature of the unique computational capabilities of structures in the cortex called cortical columns and microcolumns.

Decision Theory

Agents make decisions. The theory behind decision making has been very thoroughly explored in areas such as management theory and game theory. The agents in this case are, of course, human beings. Any information processing subsystem that affects the actions of a whole system is an agent, so the principles of decision making (and taking) are applicable in all complex systems. In this section we examine some of the general principles of decision theory as it applies to managing system behaviors. A core set of principles are applicable to all agents from the mere mechanical controllers (e.g. thermostats) to complex cyber-physical systems such as mobile, autonomous robots, to living cells, to human brains.

The structure of a decision process can be represented in the form of an expanding tree (in mathematics the term tree is given to an inverted branching structure – see Figure 10.2) in which each level represents possible states of the world that will be used to compute a decision according to the decision model employed. Each branch to a lower level of the tree represents an alternative choice based on the current state and the preferred future state that should obtain from the choice of that branch. The mechanics are that at each node in the tree the agent evaluates the state of the world (as in Figure 10.1), uses the decision model to obtain a choice of which future state (branch) should be the most desirable given the current conditions, and the actions of the agent should lead to that child node. Decision trees capture the overall structure if the states of the world are knowable and predictable based on the current state of the world, the degree to which the agent can ascertain it, and the amount of uncertainty from several different factors. A perfectly deterministic decision tree, such as a game of Tic-Tac-Toe played by two very alert players has perfectly predictable future states given the first move. If one of the players is less than alert then the outcome might be surprising.

In very complex decision processes, there are many sources of uncertainty that are not just involving mistakes made by the agent. The input data may be incomplete, noisy, or ambiguous. The decision model may be sketchy and not include all of the real possible branches (this is why learning agents evolved!) Lastly the motor outputs from a decision might not be sufficient to actuate the desired outcome in terms of the changes made to the state of the environment.

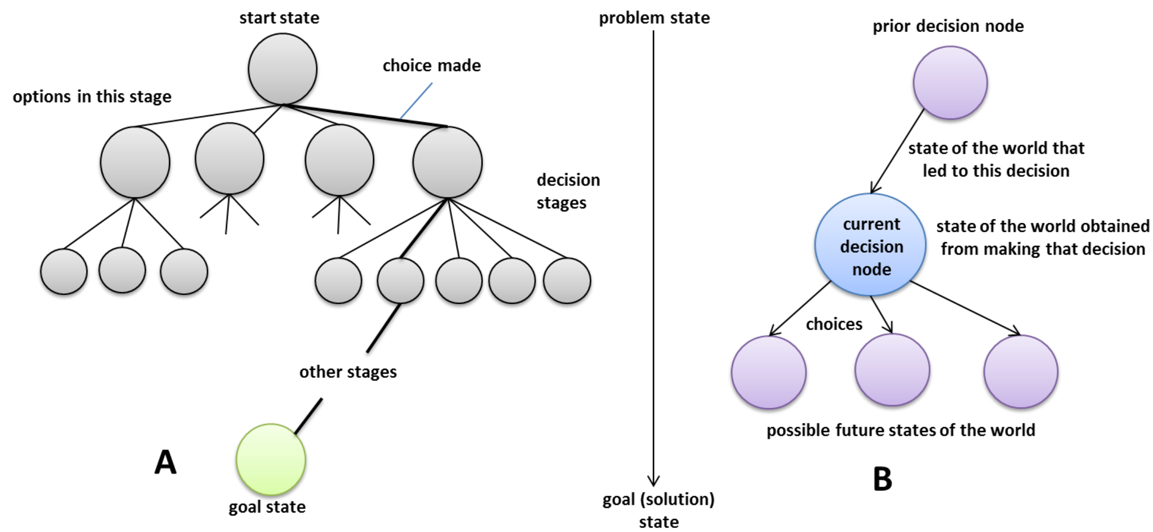

Fig. 10.2. (A) Decisions can be modelled as being made in a sequence of discrete stages, starting with the system being in the start state. The process is represented by a tree where time progresses downward. Each node in the tree represents a decision point. The child nodes of that node represent the choices that can be made to progress to the next stage. (B) The efficacy of a decision depends on how much information is available to the agent about the consequences of the next stage.

Figure 10.3 shows the situation for two decision models, one for a simple learning agent and the other for an agent having multiple sources of background knowledge that can help modulate the decision process. These are shown from the perspective of a 'current decision node' as shown in Figure 10.2.B. The lower nodes represent future possible states. In both figures the outer purple oval represents a variety of a priori knowledge that might help affect a decision. We normally think of decision making as a rational process, involving intelligence. But in animals and humans there are a number of other cognitive processes that contribute to the making of decisions (c.f. Damasio, 1994; Kahneman, 2011). Antonio Damasio (1994) was the first to alert us to the idea that when we experience the results of a decision made it usually has the effect of tagging[7] that choice (linkage in the neural representation of one state to that of the new state) with a valence – it was either a positive experience or a negative one. In the figure the small ovals overlying the links leading to the next stage contain either a plus (+), a negative (-), or an unknown valence (?). The strength of this tagging depends on the salience of the experience. Weaker valences may be altered in the future or strengthened if reinforced frequently along the lines already encoded. All other things being equal, the agent in Figure 10.3.A would tend to choose the path to choice B since at some time in the past this choice resulted in a positive

[7] Damasio's term for this tagging was 'somatic markers' in that they recorded body states or feelings, at the time the experience happened. If a person (or any animal) experienced a negative consequence vis-à-vis their body states, then the marking of the association link with a negative valence provided a kind of 'warning' that choosing that set of actions and subsequent environment state would not be immediately advantageous. See page xxx.

outcome that tagged it with a positive valence (see Mobus, 1999 for details on how this tagging actually works in neural networks).

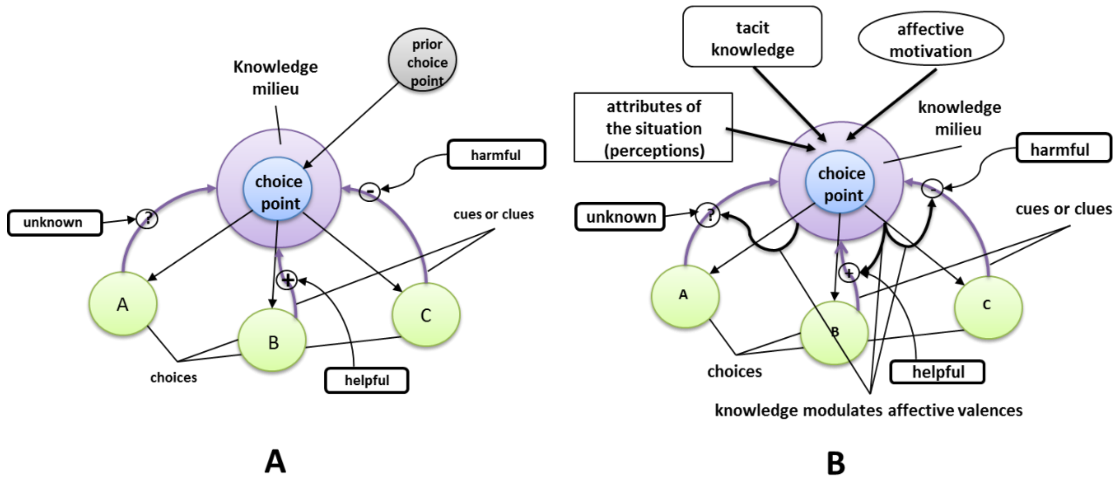

Fig. C.3. (A) The agent may learn to associate choices with various cues or clues from past experience, tagging pathways with valences (+ or -). (B) Highly autonomous agents have enriched knowledge milieu with which to judge the best choices, even if valence tags suggest otherwise.

Choices that humans make are modulated by a much more complicated knowledge milieu. This milieu is where *judgement* takes place. Seemingly rational decisions may be highly influenced by a number of factors involving memories and feelings. Humans (and likely many mammals and birds) are able to override the simple path valences if they are able to consider states much further down the tree levels. That is, it may be that taken alone a given path might be tagged negatively at the current stage but there is a path out of that node to a much more positive node further down. This has to be available in the knowledge milieu at the current choice point (Figure C.3.B). It is the reason that a chess player can sacrifice a valuable piece in a game, knowing that this will likely result in the opponent making a choice that will allow the first player to put the opponent in check (a rewarding state).

A contribution that is made in influencing decisions is the degree of confidence that the agent has regarding the choice (Burton, 2009). Both affective (emotional) and tacit knowledge seem to contribute to this factor. It is believed that confidence in one's decisions is necessary in order to take action without hesitation. It helps prevent analysis paralysis or delaying making a decision until all of the relevant data have been processed. Herbert Simon's concept of satisficing may be the effect of the sources of confidence reaching a threshold that leaves the agent feeling certain they are making a good choice. This, however is a source of error (previous citation) when the agent believes they are correct based on faulty knowledge or emotions. A sad example of this phenomenon is when those of a certain political persuasion insist that anthropogenic climate change is not happening. They may not be lying in the conventional sense, only convinced of their own veracity because their entire knowledge of the goodness of a certain economic condition forces them to conclude that global warming is a hoax.

The classical psychological study of decision making is converging with neuroscience at the point of recognizing that high level decisions are made in neural circuits (notably in the prefrontal cortex) with considerable information provided by many parts of the brain, including what we might call the more primitive parts in the core brain, like the amygdala (fear processing). As this research progresses we will be able to see more clearly the influences that emotions and biases have on our ‘semi-rational’ decision processing. Along the same lines, as our artificial agents become better decision makers with varying degrees of autonomy it will be because we have better understood the dynamics and processes.

One very important topic related to decision processing is learning and memory. In Figure C.3.B reference is made to ‘tacit’ memory. This is the realm of general mental models of how things work rather than explicit memories such as episodic (specific instances) ones. These memories are learned; mental models are constructed over time and multiple occurrences of situations involving the concept, what we call inductive learning. Categorical memory or the existence of ‘classes’ of objects, relations, and actions are in this realm. One recognizes one’s own house as a ‘house’ similar to other houses, but with particular features, including location, that makes it unique among houses. There is a tacit model of house-ness that has been constructed by an inductive process of experience over time that can then be used to model specific instances of particular houses one becomes familiar with. The particular instances are held in explicit memory accessible by conscious recall to put into working memory for thinking[8].

The amount of capacity for forming memories of mental models and the ability to learn complex causal relations and the ability to use those models in decision thinking are all factors in how much autonomy an agent might be possess. As with the example of overriding a negative valence tag on a decision path in order to achieve a better ‘payoff’ later the more capacity an agent has for ignoring affective recommendations or confidence factors the more autonomy they possess.

Degrees of Autonomy

We will classify agents by the degree of autonomy they have in making decisions. This can be based on several factors such as the capacity for modifying the experiential memory or even the decision model, but also on the relative *looseness* of the decision model itself. Autonomy is linked with the capacity for judgement modulation on decisions as discussed above. A thermostat has no autonomy. Its decision model is fixed for all time. Similarly, homeostatic mechanisms in

[8] It is generally believed that bringing a thought (concept or model) into working memory is based on activating the kernel of the concept which, in turn, activates all of its components in associative and even perceptual memory areas. When all of the components of the concept are activated and in synchrony the conscious-processing brain is aware of the object. See also, Appendix A.

living systems may have slightly more variation with which to work, but are still destined to react to their sensed environment as designed by evolution. It isn't until we get to very complex brain structures and correlated complex behaviors, particularly cortical architectures, that we see 'experimental' responses introduced into the mix. This means that a brain exposed to a novel set of inputs may construct a novel output through 'guessing' as to the correct response and recording in its experiential memory the results. It then takes further off-line processing to modify the decision model accordingly. If the novel behavior was successful in meeting the decision objective (see below) then it will become part of the repertoire of responses used in the future. Otherwise its model will be inhibited from further use.

Computer-based control systems are just entering the realm of autonomous decision processing, though as far as we know there is no model of artificial judgement in the way that we think of artificial intelligence. The success of chess-playing programs like Deep Thought[9] created an artificial knowledge milieu based on the judgements exercised by human experts, collected in a large set of heuristics that acted as judgements in playing. For now, the knowledge milieu and judgement processing has to be hand crafted. Recent work on artificial neural network models (along with statistical learning methods), called Deep Learning, are being used to automatically construct knowledge milieus for judgement-like processing in autonomous cyber-physical systems like self-driving automobiles. The methods being employed are still fairly restricted to stationary, or semi-stationary pattern learning so the nature of what is being learned is limited to minor variations on pattern sets. The major paradigm of network learning is still based on what is called fully distributed representation, in which categories of patterns are encoded in the weights of all synaptic connections through the network. This is fundamentally no different than was being used in the late 1980s when networks were constrained in sized by the available hardware of the time. All that seems to have changed is that we now have much bigger, faster computers with much more memory. So we can tackle bigger more complex problem domains. Still this kind of learning, while useful for cyber-physical systems is a far cry from the way the mammalian brain works. And so it is not able to emulate even animal learning necessary for true judgement and complex concept formation (see Appendix A for a more in-depth account of knowledge encoding in the neocortex).

Computers have been used for decades in what are called embedded systems controls. The common thermostat on your wall, today, is such a system. These are relatively simple devices using deterministic algorithms to compute the transfer function converting the ambient temperature into a control signal to the furnace or air conditioner. The HVAC (heating, venting, and air conditioning) systems in whole large commercial buildings are much more complex, and today can involve distributed computing for local zone controls, but a centralized coordinator computer working on optimization of the whole system across these zones has to monitor and

[9] See the Wikipedia article: https://en.wikipedia.org/wiki/Deep_Thought_(chess_computer) for background on IBM's program.

communicate tuning commands to the network of computers. These are among the first examples of engineered hierarchical (layered) control systems, still simple enough to be engineered for determinacy. The distributed computers have a minimum amount of autonomy relatively speaking, their agency is highly constrained.

More complex systems have been recently coming under the aegis of engineering, recognizing that increasing complexity of systems, reflecting the complexity of their operating environments, also may involve increasing their autonomy and agency. Complex cyber-physical systems, the Internet of Things (IoT), and other complex engineered systems are referred to as "systems of systems" (SoS) and the emerging approach to engineering such systems is becoming cognizant of the need to deal with increasing stochasticity (surprise) in terms of the systems fulfilling their intended purposes. More autonomy must be given to component subsystems so that they can continue to operate nominally in the face of disruptions in other parts of the embedding system. The design of the Internet is a very good example. Nodes, such as computers, routers, and switches (along with domain name servers, and other service providers) are designed to continue operations even when other components go down. The hierarchical control aspects of the Internet are handled by a distribution of roles to special nodes in the network (e.g. service providers and backbone providers) that issue global commands to the distributed operational nodes[10].

Mobile, autonomous robots are another example. Some discretion in choosing appropriate responses to unpredicted environments have to be given to the robots so that they can avoid being thwarted by those environments. Research on how, exactly, to implement this capacity is quite active. The desire is to create robots that can act as 'slaves' to human needs, recognizing that they need some flexibility, i.e. autonomy, so as to negotiate complex unpredictable circumstances (which would make them useful.) At the same time, they cannot be given too much autonomy. We would not like to implement the scenarios depicted in movies like "The Terminator" or "The Matrix."

The issue of autonomy is especially cogent in the discussion of governance and management of human organizations and state governments where we are talking about the most autonomous agents we know of, human beings[11]. These are the key decision agents operating in the structures of management and governing. The environments with which they must interact are extremely complex and full of uncertainties, requiring a high degree of autonomy and the creativity needed to deal with those environments. On the other hand, humans are notoriously guilty of constructing and holding beliefs in decision models that are not necessarily valid in terms of how

[10] We take up the concept of distributed specialists more fully in appendices D and E.

[11] It is important to note that supra-biological entities such as corporations and states seem to exhibit degrees of autonomy themselves. However, we might not be able to easily differentiate between the autonomy of individual human agents, like the managers or lawmakers, and the autonomous behavior of the whole system. This will be an area of active research in the future.

the real world actually works. Put bluntly, human beings may entertain ideologies that may seem logical but do not actually correspond with reality. At this writing a large contingent of so-called conservatives (political) still reject the overwhelming evidence for anthropogenic climate change due to the burning of fossil fuels. It has become abundantly clear, empirically, that climate change is underway and causing increasing variations in weather patterns, the extremes of which are damaging. Their disinclination to understand what is happening is actually based on a *true* proposition, that wealth production is directly tied to energy use, which has been due to the burning of fossil fuels. As we pointed out in Chapter 8, to reduce carbon emissions means reducing the burning and consequently the energy available to do economic work, in lieu of replacement by so-called renewable sources, a possibility still being investigated scientifically. Their philosophy includes the belief that economic growth is the only good for humanity. But this flies in the face of the problem presented. Thus, their only real position has to be to deny climate change (or the human contribution) in order to protect their cherished belief in wealth production and expansion.

Such ideological positions clearly compromise the effectiveness of their agency when it comes to making legitimate and helpful decisions. So, the irony of the human position as agents in decision processes is that they enjoy the maximum degree of autonomy but also the maximum degree of choice to ignore reality's messages. Human beings appear to be free to choose what they are going to attempt to learn – a meta-decision process. And they can choose not to attend to sources of factual information if there is a hint that it would require them to learn something contrary to their beliefs.

In the realm of management this had tended to be less of a problem prior to the emergence of the neoclassical, neoliberal capitalist ideology for the economy. With the latter came a heavy emphasis on maximization of profit-making and shareholder wealth as opposed to customer benefits (e.g. quality, as low a price as possible to make a fair profit, etc.), employee wellbeing, or that of the environment. This was coupled with a tremendous increase in reliance on debt financing and the financialization (making profits off of second and higher-order financial instruments treated as assets). Under this environment managers have been faced with different objectives in their decision models that have severely complicated and, one might say, muddied the experiential memories used to make decisions.

Thus, the current situation with respect to the degree of autonomy for human decision agents is problematic. Further complicating the matter is the kind of experiential memory constructs held by many human beings. Education, it can be argued, has not provided the majority of humans with the kind of broad grasp of how real systems actually work or of the implications of a systems perspective. The result is a poorly educated populace when it comes to insights into what constitutes good decisions in running everything from families to multinational corporations and states. A significant amount of research is needed to answer some fundamental questions about the efficacy of human agents in the governance and management of organizations (which is quite an understatement.)

Implications for systems analysis vary depending on the nature of the whole system with respect to its governance architecture. In the hierarchical cybernetic system described below, agents occupy relatively clear functional positions in the hierarchy (see the next section on decision types). The key guide for analyzing a system's governance and management system is to recognize the degree of autonomy *needed* by the agent at any decision node in the structure. One approach to this would be to have a good characterization of the environment of the system with respect to the degree of stochasticity in input flows and output acceptance by entities as sources and sinks. Additionally, internal subsystems may be subject to disruption from the activities of other subsystems with which they interact. Finally, systems are always subject to their own internal components failing or degrading so as to affect their performance. Analysis of these sources of uncertainty should look for the amount of stochasticity (variance in the range of the signals, etc.) and the kind (stationary, homogeneous non-stationary, or non-homogeneous non-stationary). All of these factors need to be taken into account in order to assess the amount of autonomy needed by the agents. In situations of low variability and stationary stochastics, the agent needs very little autonomy (e.g. a situation reflected in homeostasis). At the other extreme, high variability and non-homogeneous non-stationary, where predictability is minimal, a large amount of autonomy is needed. This means the decision model needs to be not only flexible but modifiable (learning) and the experiential memory facility must be extremely large to learn the larger number of possible causal relations needed to formulate a knowledge milieu.

Ideal Agents

The role of agents as decision makers in governance and management can be idealized for the purpose of establishing a basic model. Idealization serves the purpose of understanding how a governance and management system can work optimally (a concept that will be further explored in the next two appendices). An ideal agent will make the best possible decision under whatever conditions of uncertainty, ambiguity, or time constraints might prevail. Note that this is different from an ideal set of information inputs which are subject to environmental variations not under the control of the agent as a receiver. Herbert Simon (1957, 1998) coined the term 'satisficing' (a combination of satisfactory and sufficiency) to describe the heuristic process (as opposed to algorithmic) for making a sufficiently satisfactory (near optimal) decision given constraints on computation power and time. Simon recognized that humans, in particular, are not capable of being completely rational decision makers owing to these constraints (also see Kahneman, 2011 for a more detailed discussion of human decision making). He described the condition as 'bounded rationality' (1957). Under these conditions, humans try to make the best decision they can but recognize that the decisions will not necessarily be the absolute best, hoping that they will be good enough for the situation at hand. This works only if there is some leeway in the consequences that would not seriously disrupt the overall performance of the system.

Bounded rationality also applies to the probabilistic nature of the information input. There are always uncertainties about whether the messages being received are the 'right' ones and that the information content of the messages is ‘right’, i.e. free of noise. Thus, even ideal agents are subject to errors in decision making just due to the inherent problems of computation power, time limitations, and uncertainties.

All of these factors apply as well to automated decision agents operating in CASs and CAESs. But for the simplest embedded system controls, where we have lately been using very powerful computers and very large memories (think smart phone technology), complex control systems for, say, large airplanes or nuclear power plants, are still faced with uncertainties and time constraints that can cause problems in making always correct decisions. Engineering real-time controls requires extreme care. Task performance in these systems (e.g. mission or life-critical systems) must be guaranteed to complete within a specified time interval in order to prevent damage to the system (not to mention passengers riding the system!)

Human decision makers, as probably comes as no surprise, are far from ideal, even in terms of satisficing. With humans, there are several additional influences at play on decision making.

In addition to the constraints and disruptions on information input, human memories can be notoriously faulty. The decision models can be tainted by ideological beliefs that are non-veridical. Rational thinking (computation) is slow (Kahneman, 2011) and easily disrupted. But most of all, humans are strongly motivated by emotions and moods (Damasio, 1994). Human beings evolved decision making capabilities suited to the late Pleistocene era, for hunting and gathering, for small group social living, and for relatively less complex environmental contingencies. The built-in heuristics and biases (Kahneman, 2011) served us quite well in those conditions. But now, in a technologically embedded situation and in a huge scale social environment, these are proving more often detrimental to decision making within governance and management roles. One might reasonably argue that human beings are far from ideal as agents. Even so, it will be instructive, when examining the issues of governance for government, for example, to ask what would the ideal system look like given ideal human agents, i.e. those that are motivated by non-selfish sentiments, work with non-biased decision models for their particular decision types, and are versed in rational thinking. They may still be faced with having to reach satisfactory and sufficient decisions (near optimal, but not absolutely optimal), but they will play their roles in a manner that leads to success of the whole system being governed.

We will take this approach below to project an ideal government for a social system, based on the HCGS model. We will then, briefly, examine several aspects of current governments to show how they deviate from the ideal to the extent that they too often, under modern social conditions, fail to adequately provide the sort of long-term sustainable governance for their societies. In a future work, the author will pursue this line of analysis in hopes of finding a pathway from the current situation to a government architecture that is more in line with the HCGS architecture and makes provisions for helping human agents become maximally effective despite the above-mentioned problems.

Agent Roles and Decision Types

We now embed the model of agents into organizational structures in which decisions affecting a system's behavior are to be made. This will specifically be the subject of the next two appendices on governance and economy. Agent-based systems (those relying on agency) have to operate internally to maintain themselves and produce products needed by themselves and other systems in their environment. They have to obtain the resources they need, avoid dangers or respond to threats with appropriate behaviors. Thus, we will see that decision types can be classified into three general ones and two sub-types of one of the general ones. The general types of decisions are as follows. Operational – decisions relating to the real-time work processes, maintaining stability, quality, and quantity of products required. Coordination – decisions relating to keeping distributed, semi-autonomous but effectively cooperating work processes operating as an effective whole, i.e. as a system. There are two sub-types of coordination decisions. Logistical – decisions relating to maintaining optimal coordination between the internal work processes, including processes that obtain resources from the environment and those that export products and wastes to sinks. Tactical – decisions relating to coordinating the overall behavior of the system with the behaviors of the entities in the environment with which the system interacts. The tactical agents need to obtain information from those other entities in order to make decisions regarding the successful obtaining of resources and exporting of the system's outputs. In some CAESs we find a third general type, strategic – decisions relating to planning for future contingencies and coordinating the internal alterations of the system (learning and evolving) capacities to anticipate that future.

These decision types will be more deeply explained in the next appendix and will be further elaborated in the one following that.

Complex Environments

Complex systems evolved to be complex due to the fact that they were operating in complex environments. Complex environments generate more information that eventually becomes incorporated into the knowledge structures of the systems that are embedded in them. One theory of evolution leading to increasingly intelligent creatures (that are also inherently complex) is that intelligence evolved in response to the increases in information processing loads. Of course, this is really another version of co-evolution as a kind of arms race (see Mobus & Kalton, 2015, section 11.6, page 568). Environmental complexity arises from the pioneering of new econiches during adaptive radiation (say after a major die-off), but also from new species behaving in new ways that challenge other species to evolve new responses. This is called "coevolution" wherein there is a mutual selection pressure applied between species. In complex food webs there may be multiple participants. There are many dimensions to consider in the increase in information in evolving environments. For example, an increase in range in animals expanding their migration behaviors, extends the dimensions in space over which new behaviors are experienced by resident populations. This alone can increase the uncertainty any local species may experience as invaders penetrate an ecosystem.

Uncertainty in Interactions

In chapters 1, 2, and 3 we learned that to really understand a given system of interest one must first get a handle on understanding the environment of that system. One need not know the internal details of other entities or agents in the environment – they are modeled as sources or sinks only (which in the current context includes predators). But one does need to know what they are and what, in general, their expected behaviors will be in order to understand the inputs and outputs to/from the SOI.

Just as we are concerned with the complexity in the SOI, so should we be with the environment, but from a different perspective. There are several concerns about the complexity of the environment. For example, the entities in the environment should generally be viewed as stochastic processes. Their behavior can only generally be described statistically. But more than that the statistics of many entity behaviors need to be described themselves statistically! That is, many environmental processes are non-stationary. They can have statistical properties, such as a mean value and variance that change over longer time scales. For example, in the case of global warming, the mean surface temperature of the Earth is trending upward due to anthropogenic greenhouse gases entering the atmosphere during the industrial age. Moreover the variances around the mean temperature appear to be increasing as well. Non-stationarity is a complication to systems that seek to stabilize their operations within an environment.

In a similar vein, the uncertainty of future behaviors of entities, very many processes are now considered to display characteristics of chaotic dynamics and fractal structures. Fortunately, chaotic behavior is generally found to operate in a dynamic attractor basin that keeps it from turning into what we normally mean by the word chaos. Both chaos and fractal structures seem to be reasonably bounded in one way or another. Some philosophers think that they are at the root of beauty in the world!

Additional contributions to complexity include the sheer number of entities with which the SOI interacts, or could potentially interact with. On top of that there is the fact that the range of types of entities might be substantial – both the classes of entities and the particular representatives from any such class play a part.

Finally, all of the above contributions to complexity may be affected by entities that exist in the environment of the environment[12]. Remember that systems are always subsystems in some larger super-system, and by induction, that super-system is a subsystem of a yet larger super-system. And larger in space means longer in time. We typically consider the SOI as impacted by

[12] See Mobus and Kalton (2015), chapter 6, section 6.4.6.5 for an introduction to the concept of the propagation of influences from outside the nominal realm of a system and its immediate environment.

its immediate environment by a kind of “light cone-like” wrapper; that somehow our SOI is not affected by events that are happening out of view – outside of the super-system. But this is, of course, not a reasonable assumption for phenomena taking place in the mid-world of human-scale phenomena. As a simple (and frightening) example consider the comet collision with Earth ~65 million years ago. It may have been set on course as a result of a small perturbation in the Oort cloud that took place thousands of years before the collision! Events outside the normal sphere of our Ecos system could have a significant impact on the subsystems within. Figure C.4 shows this as a set of concentric rings around an SOI representing the realms of influence. The inner ring represents the immediate environment of the SOI with various other entities (sources and sinks) directly influencing the SOI. Out from there is a ring representing what amounts to the environment of the SOI’s environment – a super-system wherein other entities (only one shown in the figure) have causal influence over entities within the SOI’s super-system. In other words, these entities form a super-super-system that may be outside of the observable world of the SOI but has indirect influence on the SOI through interactions with the latter’s environmental entities.

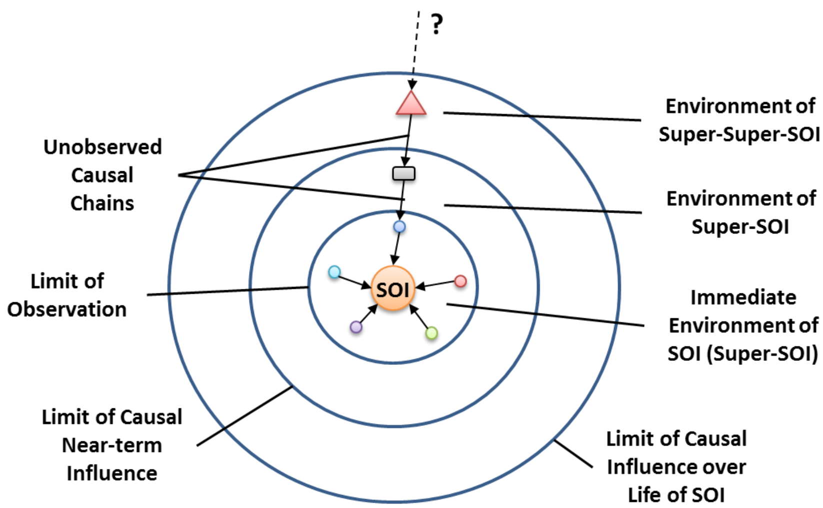

Fig. C.4. An SOI is directly influenced by elements in its immediate environment; it is the limit of direct observation by the SOI. The immediate environment, by virtue of its stronger connections with the SOI constitute a super-system that is embedded in a yet larger super-system. Due to the temporal evolution of impacts from more distant entities, and the lifetime expectancy of the SOI, current influences from the greater super-systems came from causal events in the past history of the SOI.

The same argument extends outward, as per principle #1 in Chapter 1, every system is a subsystem of a larger system until we reach the extent of the Universe. However, the further out the source of influence is found, the longer distance the influence needs to travel and so events happening to the SOI that may have been the result of changes that propagated through the causal chain from far out is also far back in time. We, on Earth, are just now receiving the photons emitted by the Alpha Centauri star cluster 4.37 years ago[13]!

[13] Source: Wikipedia article: https://en.wikipedia.org/wiki/Alpha_Centauri

The influences of distant events on the SOI could have started out being perfectly deterministic (though in general they probably are not!) But by the time they reach the SOI in some altered form they are more the source of uncertainty than determination. This time/distance tracking issue is why we cannot make perfect predictions of what is going to happen to an SOI, what the state of the environment will be at any given instant. We rely on statistical analysis to extract patterns of behavior and estimates of future behavior resulting from unpredicted events. This applies even to simple SOIs because it is a property of environments.

As a more cogent example consider the case of economic globalization being experienced today. Each nation might be thought of as a local SOI but having interactions with the larger regional and, ultimately, global economy. As this is being written one of the hottest debates coursing through political rhetoric is the negative vs. positive effects of globalization on national economies, particularly labor markets. Each local economy has become subject to interactions with the larger economies that individuals or even firms cannot directly observe, yet feel the impact[14].

Information Theoretical Considerations

From the perspective of the SOI uncertainty in interactions with the environmental entities conveys information and for CAESs results in a change in knowledge in the SOI. A change in knowledge is a change in structure that persists and leads to altered behavior of the SOI (Mobus & Kalton, Chapter 7).

A message conveys information when the receiver is uncertain about the message state (e.g. the encoding in the message) as discussed in Chapter 2. The amount of information in a message, in turn, causes alterations in the receiving system structure that allows it to dissipate future messages in that same state more readily. That is, the receiving system is much less surprised by the following messages and thus changes its own structures much less. Because of the second law of thermodynamics, the changes in structure, the knowledge, can never be perfect. And since the changes in structure involve embodied energies (e.g. covalent bonds) that can slowly 'leak' away, thus degrading the new structure, completely dissipative knowledge is never perfect.

Knowledge is stored in memory structures in more advanced CAS/CAESs. Some of these memory structures have mechanisms for maintaining long-term histories of experiences and are stable for the duration of the agent.

The amount of information, and thus the processing load on the agent, is dependent on what the channels for message input are and their capacities. Or, in other words, it depends on what the agent can attend to. An earthworm has a very limited perception of the dirt through which it crawls. The information it receives is based on only a few limited senses (mostly of a gustatory

[14] The nightly news notwithstanding. Individuals in advanced economies receive some information about their situations, but one might have difficulty verifying the correctness of the messages.

or olfactory nature – sensing chemicals). It can undoubtedly 'feel' its environment, but it does not perceive the complexity of soil. Hence it does not need much of a computing engine or experiential memory to successfully deal with its surroundings. A human soil scientist, on the other hand, sees tremendous complexities in the nature of soils. The soil scientist receives a tremendous amount of information from observations of soils. The difference is that the scientist has many more means of perceiving and receiving messages to be processed than does the earthworm. Same environment but very different amounts of information.

Computational Engine

The work processor of an agent is the computational engine that uses energy to process (transform) input messages about the state of the environment into output messages that will activate actuators to affect the future state of the environment (the double arrow from input environment to output environment in Figure C.1 is interpreted as both environments are one in the same configuration, but that the arrow from the output to the input environment implies feedback from the resulting state to the new input state after an appropriate time delay). The engine can also issue commands to change memory. In systems that can learn from experience the model may also be altered as a result of computations.

In a cyber-physical system like a mobile autonomous robot the computational engine is one or more microprocessors. In a human agent, acting on its own behalf, the computational engine is the brain[15]. For an organization, the engine is a combination of human brains and computing machines.

The characteristics of the computational engine is that it follows algorithmic or heuristic rules given by the decision model to transform the inputs into appropriate outputs (Mobus & Kalton, 2015, Chapter 8, esp. section 8.1).

Decision Model

Decision science[16] is a well-developed branch of both psychology and management science. It has been the subject of a branch of mathematics called game theory (see von Neumann & Morgenstern, 1944 for the origin) since the mid-20th century. The field is vast and fairly complex so will not be gone into in any detail here. We are concerned rather with the broad and general notion of a decision model used in the computation of decisions for action. These models, just as the computational engines, have many different instantiations in real systems. An on-line accessible introduction can be found in (Hansson, 1994).

[15] Interestingly in the human brain all three of the major components shown in Figure C.1 are combined in one fabric – the network of neuronal elements constituting primarily the cerebral cortex with some computational and communications support from the lower brain levels. In the classical Von Neumann computer model the three elements are separated as shown, though the decision model (program) is stored in a section of the general purpose memory system. The rest of the memory is given to storing data generated during the computation.

[16] See the Wikipedia article: https://en.wikipedia.org/wiki/Decision_theory. Accessed 11/27/2016.

Our major concern is with decisions made in uncertain conditions, i.e. in the face of imperfect information about the environment state due to the complexities discussed above. Since the whole purpose of a decision-making agent is to carry out actions that will preserve the system in the face of disturbances decisions need to be made in a timely fashion even when complete information is unavailable. The ideal of a decision is to find an optimal solution to a complex real-time set of requirements and constraints. The computation of such a solution using an ideal decision model in time to have an effective output is the objective. However, the constraint of completing such a computation, even with an ideal model, within a reasonable time frame is always creating a less than optimal solution. One important contribution to the issue of decision-making in uncertainty and bounded by computational complexity and time was by Herbert Simon (1969 [1996]), a theory he called 'satisficing,' in which the decision agent makes the best 'estimate' of a solution that is satisfactory for the situation. It need not be optimal, but it does need to be good enough to keep the player in the game, so to speak.

Policies - Models

Decisions are to be made against the background of the goals of the system. That is, a correct decision is qualified by how well it brings a system closer to a goal. Goals are embodied in a set of policies or rules about outcomes – what should happen as a result of the agent's decision and action.

In naturally evolved systems the policies are embedded in the structure of the system. For example, in a homeostatic mechanism, the set point represents the ideal value that a critical factor should take on. The policy is thus given in the system's construction and the procedures are embodied in the response mechanisms that act to restore the factor from a deviation.

As CAESs evolve new policies are either discovered or intentionally[17] determined (see the discussion of strategic management in the next appendix). There are two aspects of setting a new policy. The first is determining that a policy needs to be established as setting a (new) goal. The second is determining the policy mechanisms or procedures that are needed to enforce the policy.

This latter aspect becomes the decision model to be used. The procedures are built into the response mechanisms.

Real-time Decisions

The decisions that need to be made regarding the operation of a work process are made in what we generally call real-time. However, we should note that there is no one absolute time

[17] Intentionality concerns a second order decision making process in which a new goal is decided by the system, followed by the design of the mechanisms that will help the system achieve that goal. As described in Chapter 12, Design, even though a human engineer may exercise volition in choosing the specifics of the design, in many cases those choices are still basically tentative. The more novel the goal and mechanisms the more subject to tentativeness the design will be. That is, the new goal and mechanisms are exploratory initially and hence as much a product of evolution as natural mutation and natural selection are. The new goal has to be tested for efficacy.

interval that constitutes real-time. Rather, as per the mathematical definition of system given in Chapter 3 (see. Equation 3.1) every system, and subsystems, have a designated Δt that defines the absolute time interval used to measure events (state changes) at that level in the system being defined. The Δt intervals for subsystems will always be smaller intervals of absolute time. So, what we mean by real-time will depend on the natural time constants for the behavior of the system being defined as the SOI (at whatever level of analysis).

The agent(s) involved in controlling that behavior have to make decisions in relatively small increments of the Δt for the SOI. For example, the classical homeostasis controller (see Mobus & Kalton, 2015, Chapter 6, section 6.4.6.4.3 Resilience, see Figure 6.7) must respond quickly relative to the rate of change in a critical factor in order to prevent that factor from getting beyond an acceptable range of value. Examples from biology abound (e.g. blood pH regulation) but the same principle is at work in the management-by-exception[18] practice in organizations.

In all agent decision systems there are a set of inherent time lags that can affect the effectiveness of the agent's actions. There are sensing and communications time lags. These should be minimized as much as possible. Real-world sensing devices take some time to respond to changes in the values they are measuring. The output of the sensory measure then takes some time to be communicated through the channel. Electrical wires and radio frequency modulation (e.g. WiFi) are very fast modes of communications. Neural signals through action potentials propagating along an axon are faster than hormonal signaling through the blood stream. Signals reaching HSS government officials may take a very long time. Price signals propagating through the economy network (Appendix E) may be extremely slow and subject to much noise.

Other latencies are due to processing time required to reach a conclusion, take a decision. We have already alluded to this latency with respect to the time limits imposed on taking a decision discussed by Simon (1969 [1996]). The computational latency depends on multiple factors such as speed of the processor, the complexity of the decision model, and the adequacy of the sensory inputs. A second communications latency involves sending the final decision signal to the system actuator that will carry out the orders.

The actuator is also a source of latency that has to be accounted for. Real-world actuators, like sensors, require time to execute the orders. This latency depends on the power factor of the actuator – the more power applied, the faster the actuator can respond – but there are always practical limits on the amount of power that can be applied. It also depends on the inherent inertia of the system being changed.

Fundamentally, then, there is always a time lag between the time the sensor detects a difference in the value of the sensed critical factor and the time the response mechanism is deployed and effects a countering change in that factor. This is unavoidable and inherent. In the

[18] See the Wikipedia article: https://en.wikipedia.org/wiki/Management_by_exception for background. Accessed 5/31/2018.

context of the evolution or design of agents and their sensing/actuating capacities those systems that act within the real-time values of Δt for the SOI will achieve stability and fitness.

Role of Judgment in Human Decision Making

For many agents in many environmental situations the decision process can be more-or-less mechanical. If a mapping exists (either deterministic like a thermostat or non-deterministic like a Bayesian neural network (below)) then the computation is relatively straightforward. Given the state of the input messages you choose this particular output.

In Figure C.3B we introduce the role of a generalized influence on decision making in the form of what is called tacit knowledge present in the knowledge milieu surrounding decision nodes in the network (e.g. a decision tree). Tacit knowledge is the kind of knowledge that we accumulate with various life experiences that seem to have similar attributes that we integrate into our subconscious. One of the more easily understood forms of tacit knowledge is procedural knowledge or “how to do something,” which builds up over time and repetition (practice). Procedural knowledge is the basis for ‘expertise,’ the capacity to do a job (or ride a bike) without having to think about how to do it. In fact, most forms of tacit knowledge cannot be captured in language or even brought to conscious awareness.

Tacit knowledge is thought to be at work in conditioning decisions by providing background knowledge that can be applied to the current situation, which resembles the many historically experienced situations of the same sort. Perhaps, tacit knowledge is at work in what we call intuition. The later is a vague feeling that one should go one way as opposed to another even when all of the other cues and valences suggest otherwise.

Some people are better at listening to their intuitions than others. And some people have effective (good) intuitions. Such intuitions have been associated with the kind of super-intelligence we call wisdom. Wisdom is the ability to make appropriate choices in the face of uncertainty and ambiguity (both factual and moral) that will prove efficacious in the long-run (a rich literature on wisdom can be found in Steinberg, 1990a and 1990b). The problem seems to be that real wisdom is rare and often relegated to elders who have had long lives over which to accumulate veridical tacit knowledge. Thus, it seems to be the case that human agents can not necessarily be counted on to make the wisest choices as individuals. As Surowiecki (2004) points out, the wisdom of the crowds is also limited to specific kinds of problems/questions that favor a statistical estimation. Democracy (or representative forms thereof) is touted as relying on the wisdom of crowds to select the best governance agents. But experiences in recent years in a number of democratic nation-states should call this assumption into question. This issue will be raised again in the next two appendices with respect to human society governance and economy.

Decision Model Refinement – Learning

Agency

Affective Action and the Subsequent State of the Environment

Feedback Through the Environment and the Causal Loop

Effects of Complexity and Autonomy

Complexity of systems and environments is correlated with levels of organization (using Simon's hierarchy of organization concepts, 1998). Biochemistry is more complex than inorganic chemistry. Multicellular organisms are more complex than their cells. Societies are more complex than a population, and so on. As we proceed up the physical hierarchy from atoms to societies the degrees of freedom in choices increase tremendously[19]. Behaviors are more lawful at the simpler levels, somewhat rule-like at intermediate levels, and seemingly historical accident at the highest levels. So too we find that agents gain autonomy as we rise in the hierarchy. Atoms have no autonomy whereas societies are free to choose any conceivable path to behaviors not explicitly denied by the laws of physics. Human societies can choose to pollute the environment no matter the consequences, until, of course, the physical consequences are realized.

Human Beings as Agents

As noted above human beings suffer from having perhaps a bit too much autonomy, at least insofar as the constraints of reality are concerned, in the short run. They can base many decisions on guidance from ideological beliefs that do not necessarily correspond with the long-term workings of the real world.

Motivation

Purpose and Goals

[19] This reflects the amount of information that an agent must process. The more choices that an environment presents (another way of saying the number of states that a meta-system can be in) the more information is needed to process that choice.

Power to Affect

Incomplete or Erroneous Decision Models

The great bulk of human decision making depends on the nature of the mental models, both explicit and implicit, that each individual has constructed over the course of their lives. Put rather simply, two humans can have wildly different models of the same system in the world and thus different beliefs about the world and subsequent biases in their decisions regarding that decision. Perhaps the paradigm example of this is political beliefs. There seems to be a strong polarization between so-called conservative and so-called progressive views of how the governance and economics of a society should work. People from both perspectives have lived in the same world (especially the case with the advent of extensive international commerce and media) and observed the same phenomena to some degree. But both have derived different meanings from those observations. There is some growing evidence that there is a brain basis, itself based in genetic differences regulating brain development, for perceiving the world starting from these two poles. Thus, we can expect that as a person develops in life, their perceptions and consequent models of how the world works will become increasingly supportive of their innate political persuasion.

Computers and robots take over more of the work that humans have performed in the economy

Kinds of decisions left to human beings are less consequential with respect to organizational success, etc.

Decision Errors and Their Impact on the System

References

Ashby, WR (1958) Requisite Variety and its implications for the control of complex systems, *Cybernetica*, Vo1 1, No 2. Available on-line: http://pespmc1.vub.ac.be/Books/AshbyReqVar.pdf. Accessed 11/28/2016.

Beer, S (1959), *Cybernetics and Management*, English University Press, London UK.

Beer, S (1966) *Decision and Control: The Meaning of Operational Research and Management Cybernetics*, John Wiley & Sons, New York.

Beer, S (1972), *Brain of the Firm: A Development in Management Cybernetics*, Herder and Herder, New York.

Burton, RA (2009) *On Being Certain: Believing You Are Right Even When You're Not*, St. Martin's Griffin, New York.

Coen, E (2012) *Cells to Civilizations: The Principles of Change that Shapes Life*, Princeton University Press, Princeton, NJ.

Damasio, AR (1994). *Descartes' Error: Emotion, Reason, and the Human Brain*, G.P. Putnum's Sons, New York.

Findeisen, W, Bailey, FN, Brdys, M., Malinowski, K, Tatjewski, P & Wozniak, A (1980). *Control and Coordination in Hierarchical Systems*, John Wiley & Sons, New York.

Hansson, SO (1994) *Decision Theory: A Brief Introduction*, Department of Philosophy and History of Technology, Royal Institute of Technology (KTH), Stockholm. Available: https://web.archive.org/web/20060705052730/http://www.infra.kth.se/~soh/decisiontheory.pdf. Accessed: 11/27/2016.

Heylighen, F, Rosseel, E, & Demeyere, F (Eds.) (1990), *Self-steering and Cognition in Complex Systems: Toward a New Cybernetics*, Gordon and Breach Science Publishers, New York.

Kahneman, (2011). *Thinking Fast and Slow*, Farrar, Straus, and Giroux, New York.

Losos, JB (2017) *Improbable Destinies: Fate, Chance, and the Future of Evolution*, Riverhead Books, New York

Minsky, M (1986) *The Society of Mind*, Simon & Schuster, New York.

Mobus, GE (1999) "Foraging Search: Prototypical Intelligence", *Third International Conference on Computing Anticipatory Systems*, Center for Hyperincursion and Anticipation in Ordered Systems, August 9-14, 1999, Liege, Belgium, Daniel Dubois, Editor.

Mobus, GE (2000) "Adapting Robot Behavior to Nonstationary Environments: A Deeper Biologically Inspired Model of Neural Processing" *Proceedings, SPIE - International Conference: Sensor Fusion and Decentralized Control in Robotic Systems III, Biologically Inspired Robotics*, Boston, MA, Nov. 6-8, 2000, pp. 98-112

Morowitz, HJ (2002) *The Emergence of Everything: How the World Became Complex*, Oxford University Press, Oxford UK.

Meuleman, L (2008) *Public Management and the Metagovernance of Hierarchies, Networks and Markets: The Feasibility of Designing and Managing Governance Style Combinations*, Springer, New York. [“how do successful public managers design and manage combinations of hierarchical, network and market governance? In other words: what is their rationale to apply metagovernance?”]

Schwarz, E (1997). “Toward a holistic cybernetics: from science through epistemology to being”, *Cybernetics and Human Knowing*, 4, 17-50

Schwarz, E. (1992) “A Generic Model for the Emergence and Evolution of Natural Systems toward Complexity and Autonomy”, *Proceedings of the 36th Annual Meeting of the International Society of Systems Science*, Denver, Vol. II.

Simon, HA (1957) "A Behavioral Model of Rational Choice", in *Models of Man, Social and Rational: Mathematical Essays on Rational Human Behavior in a Social Setting*. Wiley, New York.

Simon, HA (1991) "Bounded Rationality and Organizational Learning", *Organization Science*, 2 (1): 125–134.

Simon, HA (1998). *The Sciences of the Artificial, Third Edition*, The MIT Press, Cambridge MA.

Sternberg, Robert J. (ed.) (1990a). *Wisdom: Its Nature, Origins, and Development*, Cambridge University Press, New York.

Sternberg, Robert J., (1990b). "Wisdom and its relations to intelligence and creativity", in Sternberg, Robert J. (ed.) (1990, pp 142 - 159).Türke, R-E (2008), *Governance - Systemic Foundation and Framework*, Springer, New York.

Surowiecki, James (2004). *The Wisdom of Crowds*, Doubleday, New York.

Varela, F (1991), *The Embodied Mind*, MIT Press, Cambridge, MA.

von Foerster, H (1974), *Cybernetics of Cybernetics*, Urbana Illinois: University of Illinois. OCLC 245683481

von Foerster, H (1981), *Observing Systems*, Intersystems Publications, Seaside, CA.

von Foerster, H (2003), *Understanding Understanding: Essays on Cybernetics and Cognition*, Springer-Verlag, New York.

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*, The Riverside Press (Houghton Mifflin Co.), New York

Wiener, N (1961) *Cybernetics: Or the Control and Communication in the Animal and the Machine* 2nd ed., Martino Fine Books, Eastford, CT

Ivenenko, VI (2009) *Decision Systems and Nonstochastic Randomness*, Springer, New York Decision system and control system are the same (chapter 2 page 9)

von Foerster, H. (2003). Cybernetics of cybernetics. In *Understanding understanding* (pp. 283-286). Springer New York

Ostrom, Elinor. 2010. "Beyond Markets and States: Polycentric Governance of Complex Economic Systems." *American Economic Review*, 100(3): 641-72.

https://en.wikipedia.org/wiki/Elinor_Ostrom

1969. The Sciences of the Artificial. MIT Press, Cambridge, Mass, 1st edition. Made the idea easy to grasp: "objects (real or symbolic) in the environment of the decision-maker influence choice as much as the intrinsic information-processing capabilities of the decision-maker"; Explained "the principles of modeling complex systems, particularly the human information-processing system that we call the mind"

- 3rd ed. in 1996, MIT Press.

https://en.wikipedia.org/wiki/Second-order_cybernetics

https://en.wikipedia.org/wiki/Management_cybernetics