

## **Can we learn how complex systems work?**

### **Abstract**

Insights gleaned from scientific analysis of complex problems risk being lost unless they are successfully communicated and understood by those making decisions. Traditionally scientists have focussed on the technical analysis of the problem, and have left it to others to ensure uptake and application of their work. Yet for many years now studies across a range of disciplines have pointed to some wide-reaching and fundamental barriers to wise decision-making in complex situations. These barriers have more to do with human cognition and psychology than the physical complexity of the problem at hand. We draw on this literature and recent preliminary studies to highlight the value of scientists paying more attention not only to the dynamical analysis but also to the human, organisational and cognitive dimensions of complex problems.

## Introduction

Is it possible for humans to manage a complex system? The belief that it is possible to meaningfully intervene in complex situations underlies our routine daily actions: we try to manage economies, ecosystems, nations, wars and our daily lives. The literature from many disciplines researching this topic presents a less clear picture. Some complex system science literature stresses the unpredictability of complex systems (e.g. deterministic chaos) yet an assumption of predictability underpins many management actions and interventions. Psychology and cognitive science characterise the limitations humans encounter in dealing with apparently simple problems, a topic we address in this paper. Artificial Intelligence has for decades attempted to devise formal, algorithmic approaches to interact with the real world with limited success (Dreyfus, 2007; McCarthy, 2007); more recently, it has directed its effort towards more flexible and adaptive approaches, often inspired by how natural systems compute and evolve, but still no one would place an economy, an ecology, a legal system or a war solely in the hands of a computer. Despite cognitive limitations, when it comes to managing natural systems we trust a human decision maker more than an automated one. Expertise takes a very long time to develop (Ericsson, 1993) and only a small percentage of people are able to develop the cognitive skills needed to address complex causal relations (Camerer et al., 1991); do all people in decision making and management positions belong to this elite? And did they all go through the supposedly needed 10,000 hours / 10 years of expertise development? What about the others? Are they unsuited to decision making?

In this paper we discuss recent literature and some experiments we have carried out or replicated. The work suggests that a) most humans struggle when dealing with the characteristic features of complex systems, b) 'experts' are not immune to such difficulties and c) these difficulties are not necessarily obvious and may be masked by the appearance of purposeful, informed decision making. We focus on the role of computer modelling in addressing some of these limitations and we discuss some initial results in this direction.

### **1 Some assumptions**

Our analysis is based on a number of assumptions and on our views on the nature of systems and of complexity. The first assumption is that managing (purposefully intervening, controlling, responding to events) requires understanding and understanding implies the ability to predict the consequences of actions. This is a contentious issue in complex system science, mostly due to disagreements on what is meant by 'prediction'. An in depth discussion in this topic is beyond the scope of this paper however several summary points are important here:

- 1) prediction means being able to anticipate limits on the expected system behaviour. For example, while it is widely known that weather forecasts are not reliable past 5-6 days, this does not mean that we have no predictability at all on the weather in 3-4 weeks time; no one would believe that the temperature in Tucson, Arizona in August could be 40°C or -40°C with equal probability. As a result no one would travel to Tucson in August with a ski jumper. The same reasoning applies to most systems for which predictability depends on time scales and resolution (Israeli and Goldenfeld, 2004);

- 2) managing implies intervening, controlling or responding to events, either to alter system behaviour according to our desire or to prepare for anticipated consequences. This, in turns, also requires some form of prediction on the likely dynamics of the event;
- 3) prediction requires some sort of system understanding, which we can describe as a (mental or formal) model used to carry out a prediction of what a system may possibly do in the future. The better the model, the better we can alter or prepare for the future (Crutchfield, 1994);
- 4) similarly, only by carrying out a prediction and checking to what extent it matches future observations can we judge whether an adopted model is effective. Modelling and predicting are two aspects of the same process.

The above premises (prediction requires anticipation, anticipation requires system understanding, system understanding requires effective models) imply that managing involves using a (mental or formal) model in order to explore possible system behaviours and their consequences and take a decision accordingly. Crucially, we also assume that managing and predicting can be learned, that models can represent systems in a manner appropriate for learning and training and perhaps more controversially, we suggest that computer models can be used to articulate the conceptualisation of a system and simulate consequences of that conceptualisation.

The second assumption we base our work on is that, *for training purposes*, it is useful to identify different features or components of a system; these features or components are not intended as subsystems, rather as contributors to the overall system complexity; for example, under this framework the presence of a feedback loop is an feature of a system behaviour which contributes to and characterises its complexity. We suggest that these features need to be understood *first* individually and *then* in combination in order to reach the global system understanding required for wise management.

| The above statement is also contentious and appears to contradict the crucial tenet of a ‘system’ approach, that is, the necessity to consider a system as a whole and to account at once for all interactions within subcomponents and between the system and the outside world. We justify our approach from a training perspective. Firstly, we suggest that for a manager facing a real-world decision, complexity manifests itself in different forms and that different skills are required for each type. We propose an approximate classification of the types of complexity and discuss what kinds of approaches are available or can be developed to address each type. Secondly, we suggest that the processes or features responsible for typical complex system *behaviours* can also be identified, studied in isolation and progressively recombined in order to train managers to detect, observe and address them at increasing level of complexity. Some initial tools we developed for this task are discussed below.

## **2 Four types of complexity**

It is a common experience that problems may appear equally complex despite being of totally different nature: to what extent the complexity of a quantum system is similar to the complexity of a social network is not obvious and using the same word to describe both may confuse rather than help any analysis. Complexity manifests

itself in different forms and for the purpose of understanding its impact on decision making we find useful to subdivide it into four types: dynamical, organisational, cognitive and behavioural or inter-relational.

In our proposed classification, the **dynamical complexity** attempts to describe the *behaviour* of the system in terms of features like system states, phase transitions, tipping points, hysteresis, oscillations, attractors and the like. This is the realm of mathematical or statistical description but the concept can be applied equally well to physical, biological, ecological and social components.

The **organisational complexity** attempts to describe the *network of interactions* which characterises a system. This also applies to the physical, ecological and social components and is represented, for example, by the political fabrics humans build to govern and by the interactions between biological species and energy sources which constitute an ecosystem. This is the realm of network analysis, ecological modelling, organisation theory and social science; it studies the paths of communication used to process information and take decisions as well as the constraints which individuals, groups and biological species place upon one another. Organisational complexity is related to dynamical complexity in two ways: the dynamic behaviour possible in a system is usually conditional on the network structure of interactions between system components; and the network of interactions itself may evolve under a set of dynamic processes.

The **cognitive complexity** attempts to describe the *challenge an individual faces* in trying to mentally organise and process the information needed to understand and interact with the system. This challenge is twofold: it includes the need to deal with an immense amount of information and at the same time with an even greater level of uncertainty. It has a contradictory flavour, since a decision maker is aware of being unable to account for all available information and at the same time wishes to have more; he/she continually struggles between the conflicting needs to simplify and to specify further. The way humans cope with this challenge and how they can be helped in this task is the topic of one of the experiments we carried out.

For behavioural or **inter-relational complexity**, we specifically mean the behavioural and psychological skills which allow some individuals to be more effective at bringing about change than others: we refer loosely to the skills which allowed Barack Obama to convince millions 'yes we can' or Al Gore to convince the world that climate change is real. We suggest this is a type of complexity, rather than merely a psychological feature, because of the intrinsic dependence on the structure, dynamics, trends and potential existing in a society (i.e. the conditions necessary for such leaders to emerge and succeed). Within complex system parlance, these are typical ingredients which lead to tipping points or phase changes in social dynamics.

It is important to notice that the above classification does not try to 'reduce' or compartmentalise a system: it does neither divide the system into physical, biological or social components nor does it divide the system spatially. Rather it corresponds to the mental challenges an observer (a stake-holder, a scientist or a decision maker) has to address in order to understand how a system works and what makes it complex *for the observer* (Crutchfield, 1994).

In the rest of the paper we describe some initial work we have done addressing dynamical, organisation and cognitive complexity; at this stage we have not carried out any analysis on inter-relational complexity.

### **3 Facing dynamical complexity**

#### **3.1 Stocks & flows**

We all know that predicting the weather is as difficult as is predicting the stock market and earthquake occurrences. These are examples of dynamical complexity, challenging even to experts. We also all know that these problems are different from predicting the fluctuations in our bank account if we know how much we earn and how much we spend, or the fluctuations in the water level of a dam if we know the rates of inflow and outflow. The latter are linear problems, which involve only addition and subtraction. We may not be good at counting but given pen and paper we expect we can sort them out; at the very least we expect experts will find such problems trivial.

In fact, this is not the case. In a series of experiments (Sweeney and Sterman, 2000; Sterman and Sweeney, 2002; 2007; Sweeney and Sterman, 2007; Sterman, 2008) Sterman and colleagues from MIT show that even individuals with high level of education (including training in mathematics) fail at 'simple' questions involving stock and flow dynamics. Here the word 'simple' has several meanings: a) it is linear, so it can be answered with high-school (primary school in fact) arithmetic b) it contains no feedbacks and c) it addresses everyday real-world scenarios. The experiments from the MIT group are exhaustive and pose the same questions dressed in different ways: filling a bath-tub, managing an interest-free bank account, storing CO<sub>2</sub> content in the atmosphere, melting ice-sheets and calculating the number of people in a supermarket. All questions can be reduced to a common simple concept: the amount of a quantity in a container depends on the difference between what is put in and what is taken out. According to the experiments well over two thirds of these mathematically proficient, highly trained, high performing individuals fail at the tests (Sterman, 2008). Higher figures occur in the general public. Sterman and colleagues make sure this is a conceptual problem and is not due to misunderstanding the questions; they ask the questions in both graphical and English form and use different representations, which do not significantly change the outcome (Cronin et al., 2009).

As naturally sceptical scientists we did not believe the results were universally true, and we assumed that experts, at least, would do better. We chose two of the published questions and ran them in exactly the same form on 4 different groups of professionals with expertise somehow related to the analysis and management of ecological systems; these groups consisted of modellers, biologists, ecologists, physicists, complex system scientists, managers, and stake-holders active in tourism development or ecological sustainability. To our amazement, the results matched the ones from Sterman and colleagues almost exactly.

Sterman and colleagues suggest that this unexpected phenomenon is due to human tendency to match patterns: typically subjects assume the dynamics of the stock

matches the one of the flow; mathematically, this means not understanding the difference between a derivative and an integral. The implications of this result are considerable: if one does not understand the accumulation process, he/she may under/over spend (creating financial damage), under/over exploit (creating environmental damage), under/over emit (creating pollution damage) etc.

Stock and flow relations are the consequence of conservation of mass, a law which all physical systems follow as do economic systems at the micro level: misunderstanding stocks and flows therefore implies that system management will fail. Said differently, understanding and *managing stock & flows is a not sufficient, but necessary, requirement to manage a complex system.*

A natural question to ask is why, if misunderstanding stocks and flows is so dangerous and yet so wide-spread, many more systems in our hands do not collapse irreversibly? There are a few possible answers. One is that the task can be delegated: a modeller may write a correct model by using the correct equations for conservation of mass and let the equations do the exact work; a manager may delegate some basic accounting to a technician who performs the exact work. Another possible view is that in fact many systems collapse for this very reason: many people run out of money, out of gas, out of water or over-pollute; to what extent this is due to a failure to understand stocks and flows may be hard to evaluate if a proper record of decision making is not available.

It is also important to ask what can be done about it. Experiments by Sterman and his colleagues suggest that a traditional technical education may not be enough, given that their subjects had studied calculus (and some of our subjects had a PhD in a scientific discipline). Our approach has been to develop simple numerical models to train subjects by letting them play with stock and flow dynamics and providing feedback on their performance. The idea is not novel, of course, and has considerable precedence in the literature. From a cognitive perspective, the purpose of the training is to use a 'correct' computer model to help a subject build an *effective* 'mental' model. When is the mental model effective? We suggest two criteria for assessing effectiveness of a mental model: when it allows the user to reliably *predict* the consequences of certain stock and flow relations, and when it becomes easier to solve a new problem. This fits nicely into our view of the relation between complexity, modelling and prediction.

Unfortunately, our experiments so far are not statistically reliable, since we worked with only 8 subjects and our results are considered to be 'anecdotal', rather than 'robust'. Within this limitation (an important one) our experiments suggest that subjects trained on computer models were able to transfer their understanding of stocks and flows to a much more complex task, represented by running a virtual chocolate factory: they controlled the balance between raw material, storage, spoilage and production better than subjects with no training in stocks and flows. Their overall performance on the management of the factory (overall economic return) was only slightly better than for the subjects with no training, which is natural since *managing stock and flows is necessary, but not sufficient to manage a complex system.* Within its limitations, principally sample size, this result is encouraging.

As for Sterman's group, in our experiment we also asked a stocks and flows question dressed in two different fashions. One question dealt with CO<sub>2</sub> and involved accumulation of pollutants (this is the same question as in (Sterman, 2008)), the other dealt with fishing and involved subtraction of resources. To our surprise the success rate on the second question was much higher than on the first. While this seems to contradict some of Sterman results, according to which the way the question is formulated does not affect the results, it suggests an avenue for enquiry, since it may suggest how a question should be formulated in order to elicit more accurate responses.

### **3.2 Feedbacks**

Feedbacks are probably the most distinguishing signature of complex systems and their dynamics may be much richer than for stocks & flows. Given the discussion in the previous section, it is unsurprising that similar cognitive difficulties are encountered by both the general public and professionals in managing a system under feedbacks. Moxnes and colleagues (Moxnes, 1998; Moxnes and Saysel, 2009) show this in a set of experiments in which subjects, including experts, overexploit a resource and as a result misjudge feedback effects. These results are particularly important for system management since overexploitation is usually explained by greed and competition and the psychology of trade-offs for the individual as in classic tragedy of the commons scenarios (Hardin, 1968). Moxnes convincingly shows how misunderstanding feedbacks may result in the same outcome (Moxnes, 2000). When this happens the consequences of misjudging dynamical complexity are misinterpreted as a consequence of social interaction, which may elicit an inappropriate policy intervention.

These considerations are confirmed by one of our informal experiments. We employed a simple conceptual model of tourism development (Casagrandi and Rinaldi, 2002) which describes the relation between environment, tourists and infrastructure and allows to study the dynamics of the system under different scenarios. The system is characterised by three different types of processes: three negative feedbacks, a positive feedback and a unidirectional impact. The subjects are asked to intervene on the system to ensure long-term benefit and they need to choose which process to target. In this particular case, the positive feedback is one that exacerbates negative impacts and so weakening the positive feedback is the key to managing the system. Nevertheless, only a small percentage of the subjects chose to target the positive feedback loop. In fact, the percentage of subjects correctly targeting the positive feedback loop is not distinguishable from random while the percentage of subjects targeting the two options which are actually detrimental for the system is higher than random; this implies that experts performed on this question at least as bad (if not worse) than would someone rolling a dice.

### **3.3 Toy models**

The difficulties subjects faced in dealing with stocks and flows and feedbacks should not be confused with what we defined as cognitive complexity. In Section 2 we defined as cognitive complexity the challenge posed by dealing with an

overwhelming amount of information and uncertainty. Neither of these features was present in the experiments we discussed in the previous two sections: subjects had all the information needed to answer correctly and the amount of information was easily manageable.

Nevertheless, there is a sense according to which the poor results on the experiments can not be ascribed merely to conceptual difficulties. It is reasonable to believe (though it may be useful to test) that most subjects would have succeeded at the stocks and flows questions had they been presented to them as ‘how does the amount of quantity in a container depend on the difference between what is put in and what is taken out?’. Similarly, it is reasonable to believe that they would have succeeded at the feedback question had they been pointed out that only one of the feedbacks was positive. This suggests that the subjects had the necessary knowledge to answer correctly, but failed to apply it; or, said differently, failed to recognise the nature of the questions or to cast it within the framework of their knowledge.

We suggest that simple models of dynamical processes can be useful as training tools to address this problem, by helping clarify the relation between such knowledge and its practical implication. Of course this idea is not novel and is the core of flight-simulator-like tools available to develop different skills.

In complex system science, ecological modelling, economic modelling and the like, models are mostly understood as scenario testing or predictive tools, but there is an increasing awareness of their potential as flight-simulators (Smith, 1994; de la Mare, 2005), training users to fly in the space of management challenges and to land on the appropriate strategy. We built some simple prototypes to test their effectiveness in addressing the challenges described in the previous sections. One, available at [http://www.per.marine.csiro.au/staff/Fabio.Boschetti/netlogo/Toy\\_Models\\_html.html](http://www.per.marine.csiro.au/staff/Fabio.Boschetti/netlogo/Toy_Models_html.html), provides training on stocks & flows problems of increasing complexity, ranging from a single stock and a single flow, to 2 stocks and 3 flows, to the inclusion of feedback. The user is given a task and can exercise at will, receiving feedback on whether the task is achieved.

Another model, available at [http://www.per.marine.csiro.au/staff/Fabio.Boschetti/netlogo/Casagrandi\\_Rinaldi\\_Mastourism.html](http://www.per.marine.csiro.au/staff/Fabio.Boschetti/netlogo/Casagrandi_Rinaldi_Mastourism.html), reproduces the tourism dynamics model described above and allows the user to verify how intervening on different type of feedbacks impacts the system. Because this system has very rich dynamics, several system responses can be simulated and understood with this tool.

With some effort it is possible to build similar models to address other dynamical processes characterising complex systems, like phase transitions, tipping points, hysteresis, and oscillations. These could stand for exercise machines in a virtual gym for system managers.

## **4 Facing organisational complexity**

As described above, the organisational complexity attempts to describe the network of interactions which characterises a system. To a certain extent we are all familiar with

this concept: when we choose who to invite to a party we account not only for who we know and we like, but also for whom the invitees know and like; will conflict arise during the party? Will someone not invited find out about the party via hidden links? Will someone dominate the party or divide the party into groups? A group of difficult friends may make party organising very complex. The art of understanding these relations is at the core of political and business success. Thanks to the dramatic rise social networking tools like FaceBook and Twitter, the concept of ‘six degrees of separation’ and the crucial role of human social networks are now well appreciated.

Often the impact of organisational complexity in complex decision making has a more abstract and subtle role and in some cases detecting the interactions among the components of a problem is a skill which requires both knowledge and intuition. For example, a modeller chooses which processes to include in a model and a software engineer decides which software modules communicate with each other and how; this demonstrates that the identification of which components and which relations are relevant to a problem (and which can be disregarded) is an essential step in both problem description and problem solving. These skills can be learnt and to a certain extent understanding feedback and stock and flows, as discussed in the previous sections, is a first step in that direction. A different body of literature discusses the educational potential of training the users in defining the system, rather than providing it a priori (Druckman and Ebner, 2008).

Yet another type of challenge lies in imagining the impact of organisation on system behaviour, all other factors being the same. We tested this idea by confronting subjects with a social dilemma. We modified a puzzle first described in Hofstadter, (1985). We asked each subject to imagine being in a live TV game show with other participants. Each player is isolated in a cubicle with a button for a few minutes. If just one participant pushes his/her button then he/she will receive \$1 million. If nobody or if more than one player pushes his/her button, no prize is given (we call this scenario 1). The question can be changed, by assuming that if no one pushes the button the prize is shared (scenario 2). The participants can not communicate during the game and we asked them what they should do and why (subjects were exposed to either scenario 1 or 2). We then asked participants to imagine the same situation with one important change: all participants get together and talk for a few minutes before being isolated and having to decide whether to push the button or not (scenario 3). The purpose of this test was to see whether the subjects recognised that in all scenarios the relation among the players is the same: since they all share both the same information and the same likelihood of winning the prize (and thus no asymmetry is present among the participants), the decision process should be the same for each player in all three scenarios<sup>1</sup>. Almost all subjects recognised this in scenario 3 and proposed to agree on a procedure which would lead to share the prize. Most subjects recognised this in scenario 2 and claimed they would volunteer not to push the button hoping the other participants would do the same and they would share the prize. No one recognised the situation in scenario 1 which resulted in most subject claiming they would push the button (quite likely preventing everyone from winning the prize) and very few claiming they would not push the button, in order to sacrifice their win for someone else’s benefit.

---

<sup>1</sup> Scenario 1 is quite challenging to analyse: the ‘optimal’ choice should be for each player to push the button with probability equal  $1/n$ , where  $n$  is the number of players (Hofstadter, 1985).

We acknowledge that the ‘optimal’ solution in scenario 1 requires a fairly deep insight which may not be available to all subjects and consequently a different test should be designed and evaluated. In the current design, the test seems to suggest that most subjects recognised the nature of the relation between the players and the consequences of an individual’s decision on the system outcome and that they choose to act cooperatively when an option was clearly available, while they chose a selfish behaviour where options for collaboration were not apparent. More experiments of this kind may give insights into the awareness of the consequences of individual decisions as a function of group organisations and help developing an intuition for detecting such situation in the real world.

## **5 Facing cognitive complexity**

Dietrich Dörner used simulated ‘microworlds’ to investigate human behaviour in complex decision-making situations (Dörner, 1996). In these experiments participants were given problems to solve in which time delays, unexpected events and counter-intuitive chains of cause-and-effect were present. For example, participants would be presented with the troubles of a fictitious village in which inhabitants have labour-intensive agricultural livelihoods complicated by human and livestock diseases, water shortages, high infant mortality and unpredictable events (e.g. drought periods). The participants were asked to make interventions in these systems with the aim of improving conditions for the simulated village inhabitants. Rather than focussing on their technical skills in solving the problems, Dörner investigated the cognitive processes and behavioural attributes and habits each participant brought to these problems. He drew particular attention to the way participants formulated their goals, the extent to which they were able to articulate the expected results of their decisions and whether they checked the realised consequences against their expectations. He focussed also on their emotional responses; for example, were failures greeted with humility and curiosity, or anger and blame-shifting?

Dörner’s work made it clear that the individual attributes that a person brings to these situations have a significant influence on their ability to make a useful contribution in a complex situation. Dörner identified a set of tangible, constructive means to improve problem-solving in complex settings. Interestingly, some of the behaviours identified by him can be in direct opposition to behaviours rewarded in high-level political and management roles.

For example, Dörner suggests that an ability to tolerate high levels of uncertainty is highly desirable in complex settings, yet the pressure on politicians and decision-makers is to remove uncertainty. Politicians find themselves making election ‘promises’. They are castigated and denigrated when later on in office these promises are ‘broken’, yet the political landscape is such that the comforting certainty of an election ‘promise’ is a myth. Strong selection pressures at work mean that there is unspoken knowledge among politicians and voters alike that promises are uncertain, yet the political discourse persists as if it operates in a realm of certainty.

As another example, Dörner makes the argument that individuals are far more effective in complex decision-making if they are willing or allowed to acknowledge

when they make mistakes and treat such mistakes as a valuable opportunities to learn. Again, powerful selection processes in workplaces and society at large inadvertently lead to mistakes being handled in less helpful ways. In some settings mistakes are masked, downplayed or ignored; it is not expected that individuals will purposefully draw attention to their mistakes when making a job application or applying for a political office, for example. At the other extreme, in the public media mistakes can be exaggerated and perpetrators publicly humiliated, with little or no opportunity for mistakes to be turned into useful learning opportunities. According to Dörner's work, these behaviours are unhelpful when it comes to negotiating complex problems.

A particularly challenging insight comes from a more abstract problem: participants are asked to manipulate arrays of lights to match particular patterns using a set of controls which have unknown and complicated effects on the light array. Participants were split into two groups: one group was asked to write down at each move their hypothesis about the effect each control has on the lights, while the other groups was simply asked to "think about their thinking". In other words, simply reflect upon their own thought processes. The self reflection group consistently and significantly outperformed the other group. Dörner concluded that 'thinking about our own thinking – without any kind of instruction – can make us better problem solvers'.

Despite good empirical and logical grounds for Dörner's identified helpful behaviours to be actively taught, practiced and fostered, remarkably little attention is given to these behavioural attributes in professional training, particularly in more technical fields of study such as science and engineering. These personal behavioural attributes were not a focus of our study, however we did request participants to complete a self-evaluation questionnaire and a personality test seeking responses which would provide information on these cognitive aspects.

In the self-evaluation questionnaire we asked about the subjects' aims and strategies their assessment of their own performance, any changes they noticed during the session and the nature of any obstacles which they confronted when playing the Chocolate Factory game. In the personality test, we employed a method proposed by Ackerman (Ackerman, 1996) which weights a subject personality according to four main traits: social, artistic, traditional and logical. The purpose of the self-evaluation and personality questionnaires was to see whether any obvious correlation could be found between learning and performance on the models versus *a priori* personality type and attitude in terms of level of confidence.

As discussed above, our results do not carry statistical significance due to the low sample size, but some interesting trends are suggested:

- 1) the personality traits, as detected by the Ackerman test, predict reasonably well the subject's performance on both the stocks & flows model and the more complex task of managing the chocolate factory.
- 2) the personality trait's ability to predict performances on the chocolate factory task is less pronounced after training on the toy models; this suggests that the toy models training resulted in some learning which partly compensated for possible *a priori* personality difference.
- 3) No correlation was found between self-confidence, self-esteem and self-evaluation on one side and performance on another: the subjects' expectations

of their own performance before the task and the subjects' evaluation of their own performance after the task did not show correlation with the actual performance. This also seems to be at odds with the stereotypical expectation that self-confidence is a desired trait for leadership.

## **6 Discussion**

In a world of ever expanding demands on limited resources, wise decision making in complex situations is crucial not only for economic and other managerial purposes but also for the welfare of humans and the environment. This is true not only for the few in positions of power but also for the general public who have to decide what initiatives to support and how to do so.

Scientific research can contribute to understanding how such decision making is carried out, how effective it is and what factors influence it. This contribution can be at least as important as more (apparently) sophisticated research in complex modelling of physical and ecological processes. It is so for at least two reasons: first, because the earth system has reached a state in which human impact is now a main driver for global environmental change (Rockström et al., 2009); second because very sophisticated scientific understanding and planning can be made irrelevant by misjudging simple crucial processes at pivotal decision making moments. The evidence we discussed in this paper, according to which even experts and scientists are prone to gross mistakes in simple qualitative judgement of dynamical processes, only adds to this concern.

Our work and the body of literature we discussed suggest two areas for further research. First, computer models can be designed to train individuals to better understand the basic processes at the core of complex systems. These models resemble flight simulators in the purpose and can be designed to cover a variety of scenarios of real world significance for decision making, including management of limited resources, unexpected feedbacks and social dilemmas. Our initial results with this approach are encouraging and indicate that it is worth developing a suite of training models covering a wide range of scenarios which are known to present cognitive challenges to human understanding, to test them more fully.

It is also reasonable to question the currently accepted view of what makes an effective manager or an effective decision maker. Introspection, self-criticism, ability to tolerate uncertainty, acceptance of own mistakes and willingness to learn from them, curiosity to unravel causal relation and ask 'why' questions rather than aiming straight at 'what' actions to take, patience in searching for evidence and counter evidence, are not features which are stereotypically searched for in leaders. However the literature, and more modestly our initial results, seem to indicate that these are essential for addressing complex questions. More widespread awareness about what makes an effective decision maker, possibly leading to improvements in training programs, may have an immense impact on a wide variety of real world issues.

**Acknowledgments:**

We thank Bill De La Mere for instigating this work by asking questions on the role of models in natural resource management. We also thank Geoff Syme who provided valuable guidance and discussion and also enabled the collaboration between CSIRO and Edith Cowan University researchers. Peta Dzidic provided helpful advice when we were developing the models and questionnaires. We were grateful for the opportunity to work with Edith Cowan University students and thank them for their participation.

## References:

- Ackerman, P.L., 1996. A theory of adult intellectual development: Process, personality, interests, and knowledge. *Intelligence*, 22:227-257.
- Camerer, C.F., Johnson, E.F., Ericsson, K.A. and Smith, J., 1991. The process-performance paradox in expert judgment: {How} can the experts know so much and predict so badly? , Towards a general theory of expertise: {Prospects} and limits. Cambridge University Press.
- Casagrandi, R. and Rinaldi, S., 2002. A theoretical approach to tourism sustainability. *Ecology and Society*, 61:13.
- Cronin, M.A., Gonzalez, C. and Sterman, J.D., 2009. Why don't well-educated adults understand accumulation? A challenge to researchers, educators, and citizens. *Organizational Behavior and Human Decision Processes.*, 108:116-130.
- Crutchfield, J.P., 1994. The Calculi of Emergence: Computation, Dynamics, and Induction. *Physica D*, 75:11-54.
- de la Mare, W.K., 2005. Marine ecosystem-based management as a hierarchical control system. *Marine Policy*, 29:57-68.
- Dorner, D., 1996. *The Logic Of Failure: Recognizing And Avoiding Error In Complex Situations* Metropolitan Books, New York.
- Dreyfus, H.L., 2007. Why Heideggerian AI Failed and How Fixing it Would Require Making it More Heideggerian. *Philosophical Psychology*, 20:247 - 268.
- Druckman, D. and Ebner, N., 2008. Onstage or behind the scenes? Relative learning benefits of simulation role-play and design. *Simulation & Gaming*, 39:465-497.
- Ericsson, A., 1993. The Role of Deliberate Practice in the Acquisition of Expert Performance. *Psychological Review*, 100:363-406.
- Hardin, G., 1968. The tragedy of the commons. *Science*, 162:1243-1248.
- Hofstadter, D., 1985. *Metamagical Themas*. Basic Books, New York.
- Israeli, N. and Goldenfeld, N., 2004. Computational Irreducibility and the Predictability of Complex Physical Systems. *Physical Review Letters*, 92:074105-074101-074105-074104.
- McCarthy, J., 2007. From here to human-level AI. *Artif. Intell.*, 171:1174-1182.
- Moxnes, E., 1998. Overexploitation of renewable resources: the role of misperceptions. *Journal of Economic Behavior and Organization*, 37:107-127.
- Moxnes, E., 2000. Not only the tragedy of the commons: misperceptions of feedback and policies for sustainable development. *System Dynamics Review*, 16:325-348.
- Moxnes, E. and Saisel, A.K., 2009. Misperceptions of global climate change: information policies. *Climatic Change*, 93:15-37.
- Rockström, J., Steffen W, Noone K, Persson A, Chapin FS, Lambin EF, Lenton TM, Scheffer M, Folke C, Schellnhuber HJ, Nykvist B, de Wit CA, Hughes T, van der Leeuw S, Rodhe H, Sörlin S, Snyder PK, Costanza R, Svedin U, Falkenmark M,

Karlberg L, Corell RW, Fabry VJ, Hansen J, Walker B, Liverman D, Richardson K, P, C. and JA, F., 2009. A safe operating space for humanity. *Nature*, 461:472–475.

Smith, A.D.M., 1994. Management strategy evaluation - the light on the hill. In: D.A. Hancock (Editor), *Dynamics for Fisheries Management*. Australian Society for Fish Biology, Perth, pp. 249-253.

Sterman, J.D., 2008. Risk Communication on Climate: Mental Models and Mass Balance. *Science*, 322:532-533.

Sterman, J.D. and Sweeney, L.B., 2002. Cloudy skies: assessing public understanding of global warming. *System Dynamics Review*, 18:207-240.

Sterman, J.D. and Sweeney, L.B., 2007. Understanding public complacency about climate change: adults' mental models of climate change violate conservation of matter. *Climatic Change*, 80:213-238.

Sweeney, L.B. and Sterman, J.D., 2000. Bathtub dynamics: initial results of a systems thinking inventory. *System Dynamics Review*, 16:249-286.

Sweeney, L.B. and Sterman, J.D., 2007. Thinking about systems: student and teacher conceptions of natural and social systems. *System Dynamics Review*, 23:285-311.