

Policy, Design and Management: The *in-vivo* Laboratory for the Science of Complex System

Jeffrey Johnson

Faculty of Mathematics, Computing and Technology, The Open University
Walton Hall, Milton Keynes, MK7 6AA, United Kingdom
j.h.johnson@open.ac.uk

Abstract. Complex systems scientists cannot by themselves perform experiments on complex socio-technical systems. The best they can do is to perform experiments alongside policy makers who are constantly engaged in experiments as they design and manage the systems for which they are responsible. In this context the nature of *prediction* in the implementation of real systems is much more complicated than it is in traditional science. The *goals* identified by policymakers change through time, and this is usually managed through the design and management processes. The combination of policy and design is the opportunity – the only opportunity – for complex systems scientists to engage and to be allowed to be involved in *in-vivo* experiments in large socio-technical systems. In turn this opens up new methodological approaches and questions for the science of complex systems.

Keywords: Complex Systems, Policy, Design, Management, Prediction.

1 Introduction

The science of complex systems differs from traditional science in many respects. One of the most important is that complex systems scientists cannot perform active experiments on many of the systems they study. For example, a scientist cannot decide to build a bridge to test predictions of traffic flows, a scientist cannot implement a stay-at-home policy to investigate the consequent impact on an epidemic, a scientist cannot implement a radical energy policy to investigate its impact on climate change, and a scientist cannot declare war or take measures to keep the peace. Initiating change assumes purpose and a view on how future systems ought to be (Simon, 1969). In democracies deciding what ought to be is the prerogative of policymakers, not scientists.

Large complex socio-technical systems have strong political dimensions, and changing them often involves huge resources and complicated planning procedures with delivery through many agents. Most policies are experiments with unpredictable outcomes, but they are generally not treated as scientific experiments and not instrumented. Put simply, only policy makers have the mandate and the money to conduct experiments on large complex socio-technical systems.

J. Zhou (Ed.): Complex 2009, Part I, LNICST 4, pp. 232–241, 2009.

© ICST Institute for Computer Sciences, Social Informatics and Telecommunications Engineering 2009

Copyrighted material

Policy, design and management: the *in-vivo* laboratory for the science of complex socio-technical systems

Jeffrey Johnson

Faculty of Mathematics, Computing and Technology, The Open University, UK
j.h.johnson @ open.ac.uk

Abstract

Complex systems scientists cannot by themselves perform experiments on complex socio-technical systems. The best they can do is to perform experiments alongside policy makers who are constantly engaged in experiments as the design and manage the systems the systems for which they are responsible. In this context the nature of *prediction* in the implementation of real systems is much more complicated than it is in traditional science. The *goals* identified by policymakers change through time, and this is usually managed through the design and management processes. The combination of policy and design is the opportunity – the only opportunity – for complex systems scientists to engage and to be allowed to be involved in *in-vivo* experiments in large socio-technical systems. In turn this opens up new methodological approaches and questions for the science of complex systems.

1. Introduction

The science of complex systems differs from traditional science in many respects. One of the most important is that complex systems scientists cannot perform active experiments on many of the systems they study. For example, a scientist cannot decide to build a bridge to test predictions of traffic flows, a scientist cannot implement a stay-at-home policy to investigate the consequent impact on an epidemic, a scientist cannot implement a radical energy policy to investigate its impact on climate change, and a scientist cannot declare war or take measures to keep the peace. Initiating change assumes purpose and a view on how future systems *ought to be* (Simon, 1969). In democracies deciding what ought to be is the prerogative of policymakers, not scientists.

Large complex socio-technical systems have strong political dimensions, and changing them often involves huge resources and complicated planning procedures with delivery through many agents. Most policies are experiments with unpredictable outcomes, but they are generally not treated as scientific experiments and not instrumented. Put simply, *only policy makers have the mandate and the money to conduct experiments on large complex socio-technical systems.*

The best that experimental scientists can do is to align themselves with policy makers, ensuring that the scientific basis of policy is sound and persuading the policy makers to tolerate the collection of data for scientific purposes. This is necessary but not sufficient for scientific experiments, since the whole nature of scientific *prediction* has to be rethought for complex systems.

2. Performing experiments on complex socio-technical systems

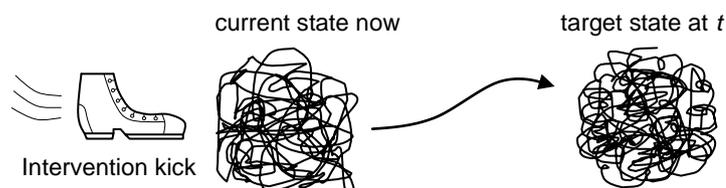


Figure 1. Experiments: predicting that given interventions will result in future system states

Complex systems scientists are seeking new methods of understanding how systems might evolve from current states to future states. *Experiment* generally involves putting a system in a particular state, making an intervention and observing the consequences. These may or may not support hypotheses on the consequences of the intervention. Since only policy makers can legitimately and practically make interventions, the best that complex systems scientists can do is make precise the desired future state(s) of the system (what policy says it *ought* to be),

make precise which system states might evolve, and suggest interventions that may best achieve the desired state(s): “tell us what the target is and we’ll tell you how to hit it”.

This is illustrated in Figure 1. Somehow the system must be *represented*. For complex systems such as a cancer or a city this representation cannot just be a few words or symbols. The representation will be complicated and generally have many levels with many relationships and numbers characterizing particular states. Characterizing the system as it is now and how it ought to be in the future involves complicated data structures holding a lot of data and computing state transitions and trajectories from these data at many scales.

The idea of a policy-oriented *prediction* is that if one makes an intervention one ‘kicks’ the system off its current trajectory onto another trajectory that will hit the target. Let s_0 be the state of system at now, time zero, and s_t be the state of the system at time t . Suppose the system is governed by *transition dynamics* with $f: s_0 \rightarrow s_t$. Let k_{NULL} be the *null kick* meaning that one does nothing to the system. Then let us write $f: (s_0, k_{NULL}) \rightarrow s_{t, NULL}$ to mean the transition dynamics of the system when no intervention is made. Let us write $f: (s_0, k_i) \rightarrow s_{0,i,t}$ to mean the state of the system at time t when intervention k_i is applied to the system in state s_0 at time t_0 .

The cooperation between scientists and policy makers assumes that the former know something about the dynamics of the system that the latter find useful. It is assumed that scientists can bring to bear better ways of representing that states of systems, s_t , their transition dynamics, f , and that this enables predictions to be made on the consequences of any particular policy kick, k_i . It also assumes that scientists are *practical*, proposing data structures that enable data collection or, more often, use of data that already exists. Also the transition dynamics must be *computable*. Policy makers are not interested in abstract predictions from scientists that do not give precise statements of what might be.

Paraphrasing the argument so far, scientists cannot do *in-vivo* experiments on most complex socio-technical systems. Generally they have neither the mandate nor the money to make interventions in large complex systems. Therefore the best they can do is sit alongside policy makers who do have the mandate and the money to make interventions. The best that a scientist can do is persuade policy makers that they know how to *predict* the outcome of any particular policy intervention, and possibly suggest particular interventions that the policy makers might find helpful.

3. Prediction and control in complex systems.

As an example of prediction, consider an ancient warship. The captain’s policy is that an enemy ship *ought* to be hit by a cannon ball. Assuming the scientist will take the commission, their job is to suggest a way of setting up the cannon in a way that can be predicted to hit the target ship when fired.

For complex socio-technical systems the reality is much more complicated than this since no-one can control all their aspects. In Figure 2 it is supposed that the original policy kick sends the system on a trajectory towards the goal, but legislation knocks it off course. For example, the policy of one Government department can be knocked off course by policy changes made by another department, *e.g.* planning regulations may become more severe, or fiscal rules may change. Figure 2 shows intervention kicks to put the system back on course towards the goal, but being knocked off trajectory by an external financial force, such as the severe global financial crisis emanating from the USA in 2007 and 2008. Again interventions attempt to put the systems back on track, and but other extreme events knock it off trajectory from the goal.

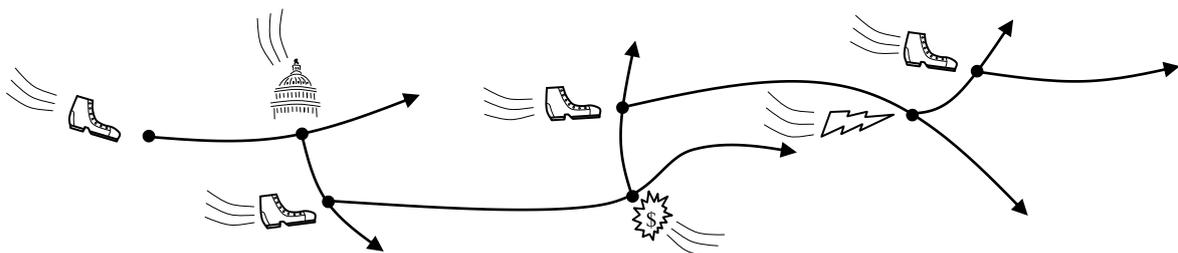


Figure 2. Policy is subject to many forces from many internal and external sources

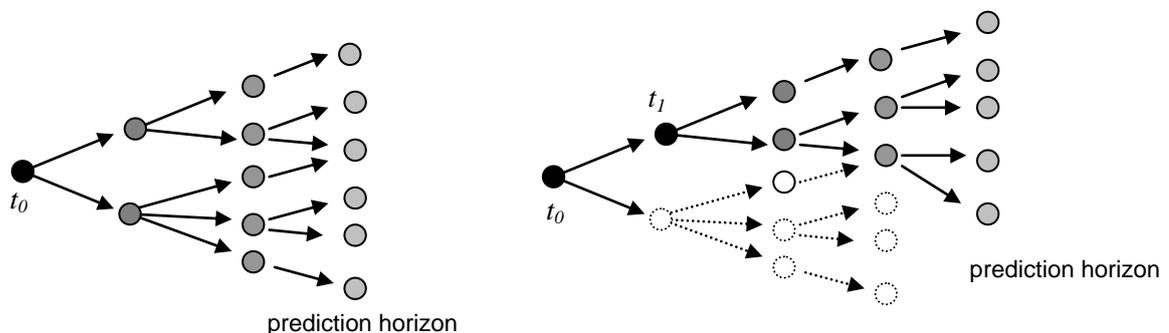
Traditional science assumes that the laboratory can be separated from its environment. Let S be the system under investigation and U the *universal system*. Then let the *environment*, E , be defined to be $E = U - S$. Thus anything not in the system is by definition in the environment. Traditional science progressed by being able to isolate S in the laboratory with E having negligible effect. For example, Galileo did not have to worry about bank rate or the weather when studying motion down an inclined plane. Often E does not have much effect on complex systems when they are behaving normally. The problem with complex systems is that they are subject to *extreme* event when they do not behave normally, and *the occurrence of extreme events is itself normal*. Understanding the nature of extreme events and the interaction and coevolution between systems and environment is part of the challenge in complex systems science.

To continue our earlier example, the cannon may be on a ship in the middle of a storm, with huge waves tossing it about. Or the regulations may change so that cannon must be fired in a different way. Or the cost of powder may change meaning that less can be used. Or a wheel may have fallen off the gun carriage. Predicting how the target might be hit becomes much more difficult as the system and its environment change.

Figure 2 is much closer to feedback control than open-ended prediction. The former is essential in real systems while the latter is almost unattainable in complex social technical systems. *Thus whole concept of prediction has to change relative to traditional science.*

4. Prediction fans and prediction horizons

Figure 2 is a simplification, since the outcome of any particular intervention kick may not be unique. In other words, an intervention may have a number of outcomes as illustrated in Figure 3. Thus the possible future states of systems fan out towards a *prediction horizon* beyond which predictions are meaningless. For example, what will the bank rate be ten years from today? Or how many cars will there be in a hundred years time? As the clock ticks the horizon moves forward, allowing some predictions to become meaningful and even useful.



(a) future trajectories fan out from the present

(b) The horizon moves forward with time

Figure 3. Predictions in complex socio-technical systems fan out and have horizons

To continue the example of the cannon and target, the analogy now is that the target ship is over the horizon and part of the prediction becomes knowing when, if ever, it will become visible and what to do until it comes in range

5. Prediction and policy in multilevel systems

The story gets even more complicated. Most complex systems have many levels, from micro-level to macro-level. When one speaks of an intervention, it may matter at what level the intervention occurs. For example, national governments can make interventions by changing laws or making funds available for their objectives. They cannot intervene at the level of the individual person, as some agencies can at a much more disaggregate level. Generally national governments do not build particular schools, this is done by intermediate level administrations. Whatever the level of an intervention, it is possible that its consequences will be felt at higher and lowest levels of aggregation.

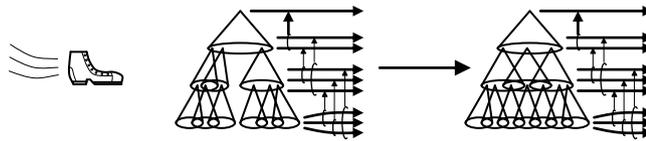


Figure 4: Predicting the consequences of a kick to a multilevel system, $k:S_t \rightarrow S_{t+\Delta t}$

To continue the analogy of the cannon, it might be that to fire the cannon requires the permission of the admiral and the captains of all the other ships, and that this permission depends on the detailed states of all the other ships and the configuration of the fleet, etc. Alongside all this are considerations of the cost of gunpowder and the lurching of the ship in the gale, and the possible position of the target beyond the horizon. No longer is the prediction the application of the theory of mechanics and the solution of a few equations calculated for a small data set. Now the prediction involves the interactions of many heterogeneous subsystems with transition dynamics dependent on large heterogeneous data sets that may be difficult to collect and very demanding to compute.

Thus every dot in Figure 3 represents the state of a multilevel system, and every arrow represents the transformation from one multilevel system state to another, as illustrated in Figure 4. Thus we have to image Figures 1 and 2 with branching, and each of the nodes in those diagrams representing complicated multilevel system representations.

Policy makers ‘kick’ their systems by making interventions at all levels and manage the consequences as best they can (Figure 4). What can complex systems scientists say about this? What does it mean to make a scientific prediction in these circumstances? Does it include the consequences at all levels?

The idea of the scientist using predictive methods to advise the policy maker how to achieve their goals is getting very complicated. But there is one final twist:

6. The goalposts move!

In the case of the cannon, it is as if the target ship changes, or the target becomes assisting rather than destroying the ship, because a more beneficial target emerges from the process of trying to hit the original target.

Consider policy makers identifying and trying to solve problems. There is the normative aspect that the system *ought* to be better, and there is the practical aspect of what could make it better. Such situations invariably have many competing dimensions. Ideally they would all be optimized, but usually the constraints compete with each other for any given solution. They have to be satisfied, meaning an acceptable compromise must be found. Sometimes the problem of creating the system that *ought* to be is over-constrained and there is no solution. For progress to be made, one or more constraint must be relaxed. Changing the constraints leads to a different problem. To be useful, any prediction made by a complex systems scientist must address the problem as it exists now, not as it was previously.

7. Design and management

By now we have entered the world of *design* which is well known to policy makers. It will be argued that the design and management of real systems is the *in-vivo* laboratory of complex systems science.

Figure 5 shows a simple diagram of the way people solve practical problems. On the left is the process of establishing what is required, which comes from those who make policy. On the right is the *generate-evaluate cycle* that characterises the design process. In this cycle designers generate new systems and evaluate them against the requirements. If the new system satisfies all the requirements it is a ‘solution’ to the design problem, and the process passes to implementation. If the proposed system does not satisfy the requirements then alternative systems are generated themselves to be tested against requirements.

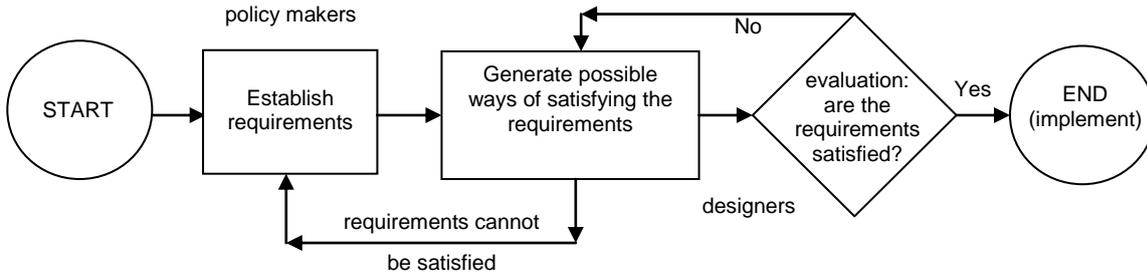


Figure 5. The simplified requirements-generate-evaluate model of the design process

The *design cycle* is a spiral process in time, since the design rarely returns to a previous state. The reason is that the designer *learns* about the problem during each iteration around the cycle. Each time a potential solution is generated the designer is making hypotheses as to what the parts of the system might and how they might fit together. Each time the potential solution is evaluated these hypotheses are reviewed and possibly tested. Each time a potential solution is rejected the designer learns something new about the system.

Anyone who has been involved in a large complicated project will recognise the process of ‘working up’ the definition of the project, and will be familiar with the interplay between constraints and possible ways forward. Policy makers know this from experience. Many scientists experience a similar process in their own projects. In this paper we are not just saying that the messy business of managing projects is necessary to administer science, we are saying that engaging in messy design and management projects *is* science, as far as complex socio-technical systems are concerned. Design and management are intimately tied up with the emerging new scientific method, with data collection, with prediction, and any possibility of testing predictions using new statistical methods.

The argument is that complex systems scientists cannot do experiments on large socio-technical systems but must align themselves with policy makers (who do experiments all the time). Policy makers execute their experiments through the design process which is a systematic way of allowing a coevolution between what *ought* to be and what *can* be.

This process involves many predictions at many levels, and the process of making predictions evolves with the requirements and proposed solutions.

8. Prediction and testing predictions *in-vivo*

Arising from the previous discussion, many kinds of predictions that can be made about complex systems. Let the experimental multilevel system be written as (B, M, H) , where B is the relational *backcloth* of the system and M is the class of mappings representing the *traffic* of system activity over the backcloth. We write $B = B_N \oplus B_{N+1} \oplus B_{N+2} \dots \oplus B_{N+L}$ where B_{N+j} is the relational backcloth of the system at level $N+j$, which can be represented by networks and hypernetworks. The mappings on the backcloth can be written $M = M_N \oplus M_{N+1} \dots \oplus M_{N+L}$ where $M_{N+i} = \{m_{N+i,j} \mid m_{N+i,j}: B_{N+i} \rightarrow Z\}$ where Z is a number system such as the real or rational numbers¹. H is a class of mappings that aggregate the mappings M_{N+i} over the backcloth, $h_{ij}: M_{N+i} \rightarrow M_{N+j}$.

Then the many kinds of predictions include:

Simple Type-I predictions: changes in mappings

Type-I-1, Fixed Level. $k: (M_{N+i}(B_{N+i}), t) \rightarrow (M_{N+i}(B_{N+i}), t + \Delta t)$

Type-I-2, Inter-Level. $k: (h_{ij}M_{N+i}(B_{N+i}), t) \rightarrow (M_{N+j}(B_{N+j}), t + \Delta t)$

¹ It can be shown that any system can be written this way, where the levels are relationally defined by parts and wholes (Type- α aggregation) and classification (Type- β aggregation) (Johnson, 2008).

Simple Type-II predictions: changes in relational backcloth

Type II-1, Fixed Level. $k: (B_{N+i}, t) \rightarrow (B_{N+i}, t + \Delta t)$

Type II-2, Inter-Level $k: (B_{N+i}, t) \rightarrow (B_{N+j}, t + \Delta t)$

These predictions correspond to a single arrow in Figure 1. They include *time series* predictions as the some of the simplest examples of Type I-1 predictions, where the future value of a mapping on a fixed backcloth is determined by its previous values. Type I-2 mappings include the aggregation of numbers over the levels of a fixed backcloth, such as the cost of building a system from components bottom-up, or the top-down distribution of resource. The employment of a new person is an example of a Type II-1 change. The impact that new person might have on the team is an example of a Type II-2 change.

Simple Stochastic Predictions

Both Type-I and Type-II predictions can be stochastic. This corresponds to putting transition probabilities on the arrows in Figure 3. For complex systems the fans may spread out very rapidly, meaning that the transition probabilities rapidly become small so that they may carry no useful information. But they may do.

Compound Predictions: multi-kick control

As Figure 2 suggests, the reality of keeping a system on a trajectory that will hit a given goal requires constant monitoring of the systems state, and constantly making new predictions on which to base new kicks to keep the system on trajectory to the target. In engineering systems feedback control is achieved by a few sensors and, usually, a few fixed equations that are used to compute changes to the control action (e.g. applying more or less power to the actuators). In such systems one does not predict that any particular control action will ensure that the system hits the target, but one can predict that the multi-kick control regime will ensure that the system will hit the target.

For complex socio-technical systems this means the empirical scientist must propose a control regime which combines Type-I and Type-II predictions across multilevel systems. To my knowledge no complex systems experiment has ever been conducted in these terms.

The reason for this is probably because complex systems are not designed and implemented in this way. The evolution of the control system occurs during the design process, which is not widely seen as part of the scientific method, even though this view is well known in design theory (e.g. Herbert Simon, Ross Ashby).

9. Design Predictions

As we have seen, design is the ultimate test of prediction since it requires an understanding of the system dynamics of the specification-design process. On the UK Embracing Complexity in Design project (Johnson, *et al*, 2007) it emerged that

- designing complex systems requires a scientific understanding of their dynamics
- design processes can be complex, e.g. manufacturing processes, supply chains
- the environment of design can be complex, e.g. regulation, fashion, economy
- design is a complex collaborative cognitive process

Complex systems scientists can contribute to the design and implementation of complex systems by providing system models allowing simple predictions to be made. To implement systems requires an understanding of the processes involved, since these play an important role in the selection of the target(s) at any time.

There are no quantitative models of complex systems prediction that take into account the environment within which systems are design, but many descriptive models of the design process.

Design as a complex collaborative process involves iterating around the coevolutionary cycles of Figure 5. Usually the 'client' occupies the left part of the diagram, deciding what they want and don't want, and what they like or don't like. Here the client is shown as policy makers –

the people that have the mandate to decide what *ought* to be and the money to commission its design and implementation.

Design is characterised by juggling constraints and making compromises, from the sketch stage all the way through to the blueprint and implementation. Even when new systems are being constructed constraints may change as new problems are discovered. The design process implicitly involves designing the *management* of the systems when it has been built.

From conception to delivery and day-to-day performance, the design process involves predictions of many kinds at many levels. This is the opportunity – the only opportunity – for complex systems scientists to engage and to be allowed to be involved in *in-vivo* experiments.

10. The new statistics

As presented here, prediction in policy and design is much more complicated than in conventional experiments, which are generally contrived to be as simple as possible. What can it mean to *test* a prediction in this context? Certainly conventional statistical techniques can be applied to local predictions, but how can the ‘correctness’ of the design process be tested in a rigorous statistical way? Of the many complications, what does it mean to make a prediction of a multilevel systems? In physics we are content that the gas laws give highly reliable predictions at the macrolevel, while the states of individuals are unknowable at the microlevel. In complex socio-technical systems the behaviour of individuals at the microlevel can have massive effects at meso and macro levels. This suggests that isolated single-level predictions will not do for complex systems, and that statistical tests will themselves have to be multilevel.

Conventional statistical methods were not developed for the kind of large complex multilevel systems discussed in this paper. For example, in the UK the major North-South M1 Motorway has been redeveloped considerably over the last few years, increasing the number of lanes in each direction to four, redesigning many intersections, and replacing a number of bridges. Such a project involves hundreds if not thousands of interacting predictions. What methods could be developed for testing any or all of those predictions? The discussion in this paper suggests that there is a completely new approach to statistical analysis waiting to be discovered and developed.

11. Conclusion

It has been argued that complex systems scientists cannot by themselves perform experiments on complex socio-technical systems, and that the best they can do is to perform experiments alongside policy makers who are constantly undertaking large and small experiments. In this context it has been shown that the nature of *prediction* in the implementation of real systems is much more complicated than it is in traditional science. In particular, the *goals* identified by policymakers change through time, and this is usually managed through the design and management processes. It is suggested that the combination of policy and design is the opportunity – the only opportunity – for complex systems scientists to be involved in *in-vivo* experiments. In turn this opens up the need for new methodological and statistical approaches for the science of complex systems.

References

- Ashby, W. R., *Cybernetics*, Chapman & Hall Ltd., London, 1956
- Johnson, J. H., Alexiou, K., Creigh-Tyte, A., Chase, S., Duffy, A., Eckert, C., Gasgoine, D., Kumar, B., Mitleton-Kelly, E., Petry, M., Qin, S-F., Robertson, A., Rzevski, G., Teymur, N., Thompson, A., Young, R., Willis, M., Zamenopoulos, T. ‘Embracing Complexity in Design’, in *Designing for the 21st Century: Interdisciplinary Questions and Insights*, Tom Inns (ed), Design Council (London), 129-149, 2007.
- Johnson, J.H., ‘Science and policy in designing complex futures’, *Futures*, **40**, 520–536, 2008
- Johnson, J.H., *Hypernetworks in the science of complex systems*, Imperial College Press (London), forthcoming 2009.
- Simon, H., *The sciences of the artificial*, MIT Press, 1969.