Systems Thinking: an introductory essay

D. Wastell, March 2012

> *Design rules! Design must be at the heart of every business… Design thinking and systems thinking are one and the same. In great design, form and function come together seamlessly. Every part contributes to the whole in a way that seems inevitable. So too in a great system.(Peters 2005)*

Systems thinking is very much in vogue, in the public services especially (Wastell, 2011). The term embraces a plethora of specific meanings, methods and affiliated sects. Flood & Jackson (1991) provide a handy overview of 8 distinctive approaches, including Soft Systems Methodology and Systems Dynamics, as well as their own 'Total Systems Intervention'. In this essay, I provide an overview of two well-known approaches, Soft Systems Methodology (SSM) and the structural modelling technique of Peter Senge which focuses on the dynamic behaviour of systems.

Soft Systems Methodology

SSM was developed by Peter Checkland of Lancaster University; it is known well in the UK but also has a strong international following. It is described copiously in a series of books written by Checkland and various co-authors (e.g. Checkland, 1981; Checkland & Scholes, 1990) and its pedigree is long. Checkland describes how SSM developed through 25 years of action research, aimed at dealing holistically with "messy real-world situations" using systems thinking (Checkland & Holwell, 2005). The empirical grounding of SSM is impressive. Checkland speaks of "several hundred systems studies carried out by Lancaster teams" (*ibid.,* p.12); the research, though based in the university, has always been conducted externally, working with "managers of all kinds and at all levels try[ing] to cope with the complexity of life's rich pageant".

Checkland makes an important distinction between 'hard' and 'soft' systems. Hard systems are typified by those dealt with by engineers, ranging across the spectrum of complexity from a domestic thermostat at one pole to a petrochemical plant at the other. Such systems can be rigorously defined and specified and can ultimately be realised as physical entities. Their design and optimisation is assisted by formal methods from mathematics and operations research; alternatives can be modelled and choices made on the basis of defined technical criteria. However, in many problematic situations faced by managers, such an engineering approach is not appropriate; indeed in such 'soft' situations it is the very inability to stipulate objectives, or even to define what the system really is, which "caused the situation to be regarded as problematic in the first place" (*ibid.,* p.12). Even situations which may seem to be simple, are not as simple as it may seem at first sight. Checkland gives the example of Concorde. Could the Concorde project simply be regarded as a system to create the first supersonic airliner? Or perhaps, as the name implies, it was 'really' a system to persuade the French that the British could be good European partners, nailing the myth of perfidious Albion. Or should it be seen as a system to develop UK aviation expertise? Of course, it was arguably all these things. What matters is that the project (i.e. the system) would be designed differently depending on the relative priorities. If the second

definition were the overwhelming priority, then there would be no ultimate need even to deliver a working product!

Spurred by these considerations, Checkland developed a new approach, based on "the fact that all real-world management problems have at least one thing in common: they contain people interested in taking purposeful action" (*ibid.,* p.13). Checkland coined the concept of the Human Activity System (HAS) to designate the web of activities linked together such that the whole set accomplishes some defined goal. Notations for producing diagrammatic representations of such 'soft systems' were developed. These can be seen as blueprints for how the world 'might be', their role being to help in problem-setting, i.e. helping understand what the problem is and aiding in its solution or amelioration. Human activity systems are not real, stresses Checkland, they are conceptual, existing only in people's heads, ways of 'making sense' of how things are and how they could be.

SSM assumes a fluid social world, one which persists and changes, "continuously socially created in never-ending social processes". Just as for Concorde, there will always be different interpretations and opinions as to what is really going on, or what should be done: as many different worlds as there are participants and spectators. To accommodate such complexity, Checkland has a brilliantly simple solution: build a model for each relevant perspective! Checkland adopts the German word *Weltanschauung* (world-view) to refer to these different orientations. It is the immanent Weltanschauung which makes:

> *…a particular model meaningful, since the purposeful action which one observer perceives as freedom fighting will be perceived as terrorism by another observer with a different taken-for-granted image of the world (*ibid.*,* p.13).

Checkland also adopts the term 'holon' from Arthur Koestler for abstract entities which are autonomous wholes and, like Russian dolls, potentially part of larger wholes. An example is urgently needed, methinks! Figure 1a shows a simple HAS. The hypothetical context (visible in Figure 1b) is higher education and the problem being addressed is the need to produce students with qualifications suited to the needs of potential employers. The notation is simplicity itself; bubbles denote activities and links indicate logical dependency, i.e. if A is dependent on B, this means that without B, A could not be done.

A set of activities is shown which certainly seem to be relevant to the issue at hand. But is this a system? It might seem so, but it is not. To see why, let me recapitulate the meaning of the ubiquitous term 'system' in the technical sense in which Checkland uses it and in which it is used more broadly in the world of systems thinking. Crudely, a system has the following characteristics: it is goal driven, purposefully transforming inputs to desired outputs, and feed-back loops are present (performance measures) to ensure that goals are achieved. Systems also operate in a defined environment, to which they are 'open'; they are encased by a permeable boundary and are capable of adaptation to changing circumstances, ensuring viability. None of the above properties is present in Figure 1a: there are no goals or feedback loops, there is no boundary and hence no environment, no means of control or adaptation. In short, no more than a 'headless chicken'. Clearly, more design work needs to be done to produce a fit-for-purpose system; this is where SSM comes in.

Like all methodologies, again we have a formal step model of how to do it (seven stages in all) but Checkland is clear that this is for expository and teaching purposes only, i.e. for the novice. Once basic competency has been acquired, the formal process becomes internalised and used in a much more flexible and fluid fashion. Checkland uses a sporting analogy:

> *…initially the thinking of the apprentices was SSM oriented. Everything that was done started from a reference to the methodology. It was noticeable that by the end of the work... [they were] using SSM more as a set of internalised guidelines which helped the attack on complex problems. The schoolboy batsman learning his craft thinks consciously about getting his left foot to the pitch of a good length ball, keeping his elbow up and swinging his bat through a vertical arc. Only when he has stopped thinking consciously about these things ... can he begin to be a real batsman (Checkland & Scholes, 1990).*

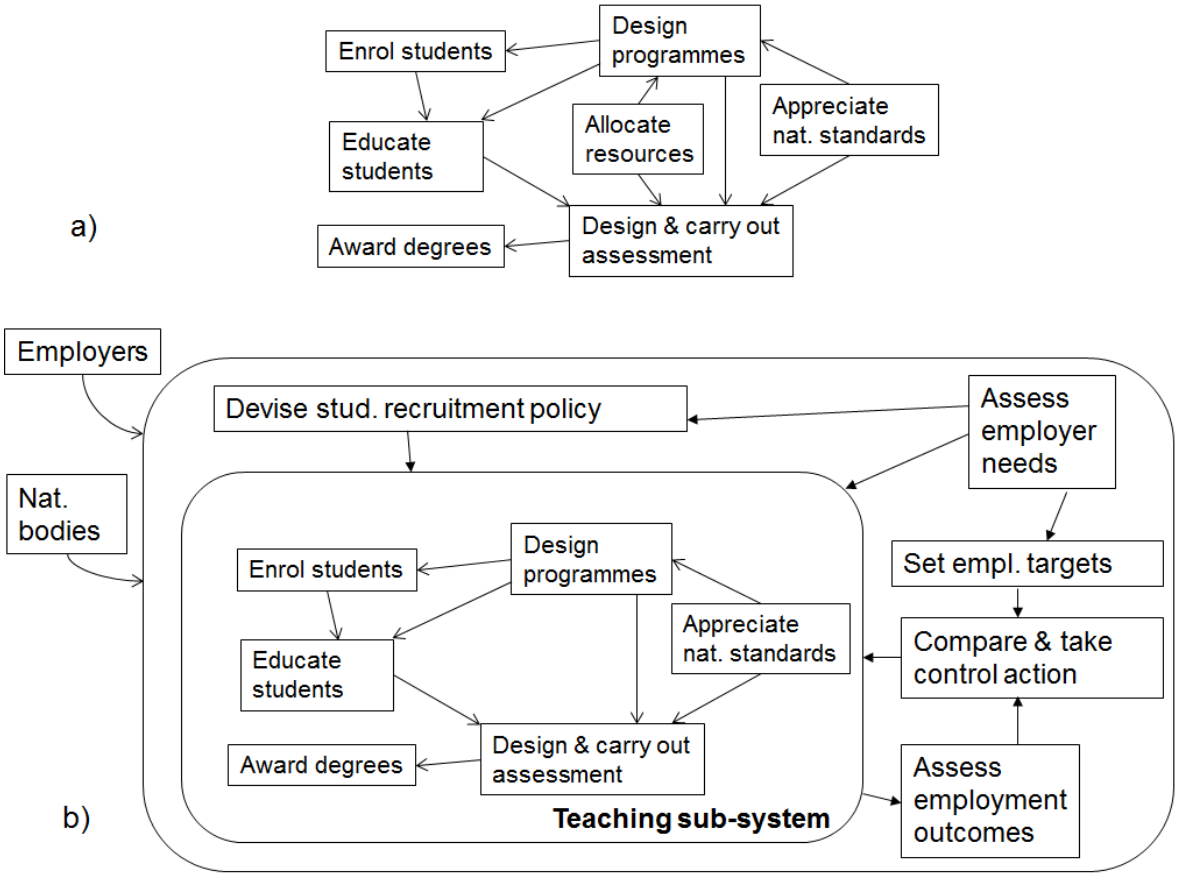

Figure 1: An example conceptual model for a university system

The seven steps of SSM are shown in Figure 2. Steps 1 and 2 take place in the 'real world'; they reflect the need for the SSM practitioner to understand fully, not the problem as such, but the problem situation, i.e. the organisational setting in which the problem is located. This 'zooming out' is critical for the reason already stated, i.e. the problem is imperfectly known; it guards against the danger of superficial diagnosis, with the real malaise going undetected. A deep understanding is thus sought in these early stages, to be articulated by drawing a 'Rich Picture', i.e. a detailed pictorial representation of the problem situation. There is no handy syntax for this: Rich Pictures are free-hand sketches or doodles and can contain as much pictorial and symbolic information as the author chooses.

An example is provided in Figure 3. It derives from a project undertaken in 2003 in a local authority IT department at the time when "e-government" was reaching its zenith. The IT department felt under considerable pressure at their time and were attempting to develop a strategy to respond to the new pressures, Ominously the wind of change is shown at the top left, in the form of central government pressure. Internally, the vicissitudes of an overloaded bureaucracy are depicted with piles of paper and jangling telephones, and e-government is portrayed as a threatening "e-train", with future plans shown in the thought bubble on the bottom right. Like Figure 3, a good rich picture typically shows up different factions in the organisation, the presence of political conflict, important physical features (crowded offices, etc.), relevant internal and external threats and possible crises. The finished picture may be useful to others, but its real value is the way it forces the creator to think deeply about the problem situation, to understand it well enough for pictorial representation. Rich pictures may be an artistic challenge for some, but the real test is anthropological. If you're stuck for what to draw, that probably means you don't understand what's going on!

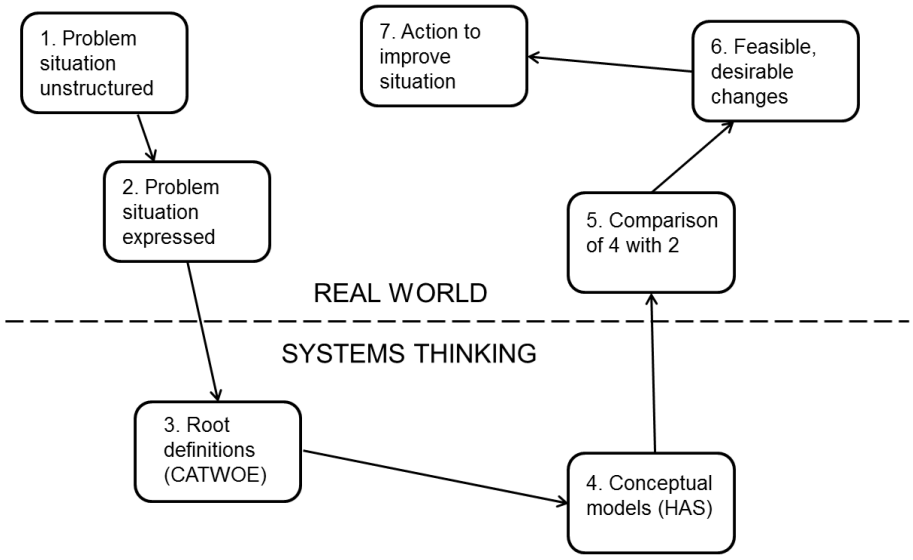

Figure 2: Overview of Soft Systems Methodology

At this point, the analyst leaves the real world and moves into the world of systems thinking (Figure 2). It is time to build conceptual systems, which will inform the problem-solving process when the analyst re-emerges, blinking, from the grotto of 'pure reason'. Elaborating a HAS is a two stage process. From her real-world research, the analyst will have identified a number of relevant but, as yet, inchoate ways of thinking about the problem situation. Let us return to our university example. We shall assume that the original problem concerned the need to increase student recruitment. One idea emerging from, for instance, the interview with the Director of Teaching, would be to increase the employability of graduates. Sitting in her systems grotto, the analyst attempts to formalise this idea, first by formulating a 'root definition' for a HAS which would achieve this function.

Six elements, represented by the so-called CATWOE mnemonic, must be specified. The essence of good design is of form following function, so let us begin with the transformation (T); systems are functional entities after all. In this toy example, the analyst decides to formulate the transformation as follows: the HAS takes candidate students (input) transforming them into marketable degree holders (output). The remaining elements of the root definition are as follows: the Customers (C) are the students, the Actors (A) are the university staff, the Owner (O) of the HAS is the university governing body, and the *Weltanschauung* or worldview (W) which makes sense of this system is the belief that awarding degrees is a good way of demonstrating the qualities of candidates to employers. The CATWOE formula also requires that relevant features of the environment (E) be addressed, for example national educational and assessment standards. Together these various elements can be combined to give the following root definition boiled down to a single sentence: "a system to award degrees (in accord with national standards) to demonstrate the capabilities of candidates to potential employers".

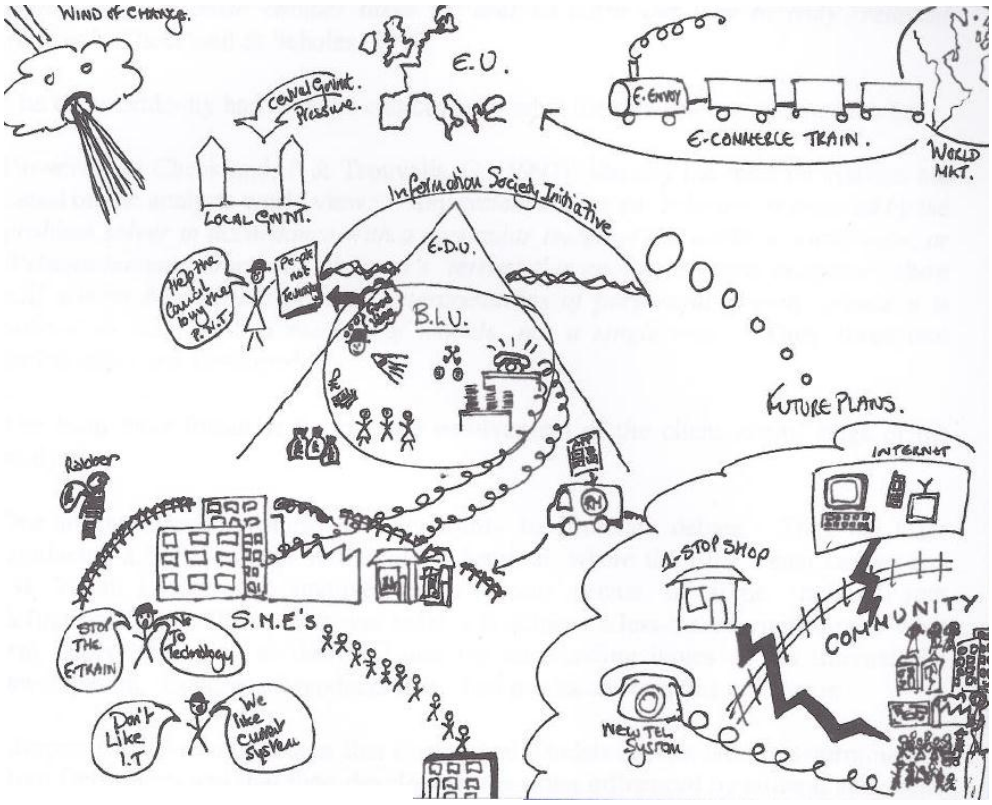

Figure 3: Example of a rich picture

A "detailed" activity model is then built in stage 4, informed by the CATWOE. The analyst's first attempt was not very good. She tries again, aiming for a "viable system" that will fit the bill. Her second attempt, Figure 2b, is much better! There are boundaries, goals and feedback loops. Above all, we see the teaching system (of Figure 2a) embedded in a management system, which determines employer needs, sets recruitment policies and sets targets for the system's success in terms of employment outcomes. Irrelevant activities have also been pruned, i.e. the resource allocation process (it is not critical to specify this, such logistics can be assumed). Now this system would have a chance of working!

At this point, we re-join reality; the method in the madness is ready to be revealed. The activity model is not the real-world, it is a hypothetical world, which can be compared with the real one, to generate ideas for action. This is the crux of SSM. With "a handful of models of this kind, of purposeful activity built from a declared point of view" (Checkland & Holwell, 2005, p.13) a coherent debate can now be held about the problem situation and what can be done to improve things. This is the so-called comparison stage of the methodology (stage 5). The different models could reflect nuanced variations around a common theme (of increasing employability), but some radically different perspectives may have materialised in the analyst's peregrinations. For instance, the university might seek to improve its recruitment by enhancing its research profile. Here quite a different HAS would be relevant, dealing with funding-body priorities, the appraisal of research strengths and those of competitors, the development and management of a research strategy, and so on.

As noted, the purpose of the debate so stimulated is to enable the organisation to think collectively about desirable and feasible changes which could be made to move things forward (step 6). In the present case, setting up a system for appraising employer needs and graduate outcomes could be one such practical option, to be implemented in the final stage. Occasionally, an overall consensus may be found but, in general, Checkland acknowledges that this debate will involve finding accommodations between divergent interests, hence "feasible" as well as "desirable". Organisational politics, like all politics, is the art of the possible. Nor is action the end of the process: rather action is where the learning really begins; as Kurt Lewin quipped, "if you truly want to understand something, try to change it!" Checkland neatly sums up the over-riding philosophy of SSM as follows:

> *SSM emerged as a learning system. In principle the learning may go on and on, and to end a systems study is to take an arbitrary step, since problematical situations will continue to evolve and will never be free of differences of interest, opinions and values (Checkland & Holwell, 2005, p.14).*

Systems thinking – Peter Senge

> *Paradox is the technique for seizing the conflicting aspects of any problem. Paradox coalesces or telescopes various facets of a complex process in a single instant.* (*Marshall McLuhan, The Book of Probes).*

In this section, I consider another influential approach within the systems thinking fold, namely the seminal work of Peter Senge on the 'Learning Organisation'. Senge heralds five "component technologies" as the vital ingredients "in building organisations that can truly learn". The first four are: *Personal mastery*: "the discipline of continually clarifying and deepening our personal vision, of focusing our energies, of developing patience, and of seeing reality objectively" (*ibid.,* p.7); *Mental models*: "deeply ingrained assumptions, generalisations that influence how we understand the world and take action…. The discipline of working with mental models starts with turning the mirror inward, learning to unearth our internal pictures of the world and hold them rigorously to scrutiny" (p.9); *Building shared vision*: "If any one idea about leadership has inspired organisations for thousands of years, it's the capacity to hold a shared picture of the future we work to create" (p.9); *Team learning*: "When teams are truly learning, not only are they producing extraordinary results but the individuals are growing more rapidly than could have occurred otherwise" (p.10).

To these four elements, Senge adds the all-important fifth, that of *systems thinking*, "the discipline that integrates the disciplines, fusing them into a coherent body" (p.12). Systems thinking is:

> *The discipline for seeing wholes…. Today we need systems thinking more than ever because we are being overwhelmed by complexity…. Systems thinking is a discipline for seeing the structures that underlie complex situations (Senge, 1990, pp.68-9).*

In common with this book's argument, Senge also sees leaders as designers *par excellence*, but acknowledges that this view of the leader as designer is not one which most managers would recognise. Senge offers the metaphor of the organisation as an ocean liner: "the neglected leadership role is the designer of the ship… It's fruitless to be the leader in an organization that is poorly designed. Isn't it interesting that so few managers think of the ship's designer when they think of the leader's role" (*ibid.,* p.341). Senge contrasts the conventional view of the leader with the more modest figure of the designer: "those who aspire to lead out of a desire for control, or to gain fame, will find little to attract them to the quiet design work of leadership", although deep satisfaction will be gained "in empowering others and being part of an organisation capable of producing results that people truly care about" (p.341). The design work of leaders includes "designing an organisation's policies, strategies and systems… integrating the five component technologies".

For Senge, what is required in the world of business is a fundamental shift in the managerial mindset (denoted by the Greek word metanoia). He proposes a number of practical techniques whereby this reorientation can be facilitated, focused on the development of each of the 5 disciplines. Here we shall be principally concerned with his concept of systems thinking, for which he proposes a simple tool called 'structural modelling'. Senge defines structure in a specific technical way, quite distinct from the conventional architectural sense in which the term is used with reference to organisations (i.e. static descriptions such as the organisation chart). In systems thinking, structure refers to the pattern of dynamic relationships amongst key components (variables) of the organisation. Senge employs the infamous 'beer game simulation' (a scenario involving actors up and a down a supply chain ordering, selling, producing, distributing and warehousing a new beer) to demonstrate how the structural properties of systems dominate over the agency of individual human actors. Even when experienced managers play the game, unstable patterns of boom and bust are produced, despite the fact that consumer demand remains stable throughout. The oscillations derive from the lags between cause and effect, the dominance of local knowledge and the failure to understand the dynamics of the system as a whole. Above all, what the beer game teaches is that structure is more important than individual behaviour: "when placed in the same system, people, however different, tend to produce similar results" (Senge, 1990, p.42).

For Senge, the primary task of 'systems analysis' is to understand such 'dynamic complexity'. This involves learning to recognise generic patterns of system behaviour ('archetypes') which recur again and again ('circles of causality'). Distinguishing two forms of feedback is critical: *reinforcing (amplifying) feedback*, when a small change builds on itself, causing a self-fuelling escalation in the original effect; and *balancing feedback*, which

operates whenever there are goals or targets that serve to regulate behaviour[1]. Whereas the former intrinsically produces instability, including exponential growth and possible catastrophe (such as the self-confirming prophecy of a bank run), the latter produces stability. Being able to identify these two forms of feedback and appreciating how they interact is the key to understanding organisational behaviour. The presence of lags (i.e. delays between causes and effects) is of central importance; it is such lags which make the management of complex systems so difficult. The so-called 'shrink-swell' effect, for instance, refers to the delay between interventions to improve performance and the realisation of demonstrable gains. An initial dip is inevitable, as the innovation is assimilated. It will always take time for new ways of working to be learned and incorporated, before the benefits are tangibly seen; the danger is that if managers do not understand such delays (or work under a management regime which makes no such allowances) then precipitous decisions may be taken, e.g. to abandon the new approach and try something else, again with the same dysfunctional effect. That way lies instability, decline and failure.

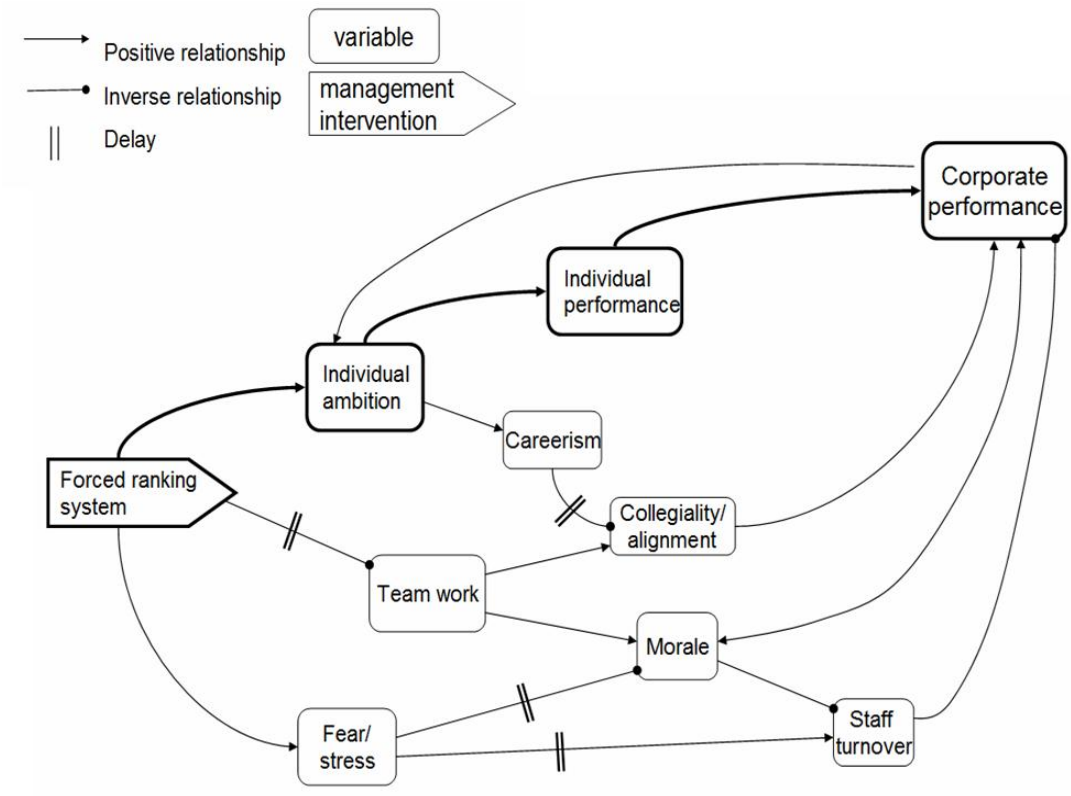

Figure 4: Illustration of a structural model

Figure 4 provides an example of such systems thinking, using my version of Senge's notation. I have used the example of forced ranking, a system of performance-related pay whereby top-performing executives are given outsize rewards, the middle 70% are "targeted

[1] Positive and negative feedback are the more usual engineering terms.

for development" and the bottom 10% are counselled or fired. In some companies, this method has seemingly succeeded and is strongly promoted by management consultancies. Yet when the evidence is interrogated, the effectiveness of forced ranking is shown to depend critically on the degree to which work relies on cooperation (Pfeffer and Sutton. Figure 4 consider the full range of the possible effects of this management intervention. Key variables are shown in boxes and the arrows linking them show the potential causal influences which can be imagined. Angular arrow heads show positive links whilst blunt terminals show inhibitory influences, where an increase in one variable (e.g. morale) produces a decrease in its dependent (e.g. staff turnover), and vice versa. Feedback loops are not explicitly shown. Senge provides some specific notation for distinguishing reinforcing and amplifying feedback; here such loops have been left implicit. They can readily be identified though. Setting targets for individual performance, for instance, provides an instance of balancing feedback, acting to stabilise the individual's behaviour at a particular level, assuming the level is within practicable reach.

Looking at the system in Figure 4 as a whole, the crunch question is: what will be the overall outcome of forced ranking? The flow emphasised in the thicker links, through individual ambition and performance to enhanced organisational performance, is what we want to happen. The rest of the reticule of causal ramifications draws attention to all of the other possible effects when the system is considered as a whole, many of which have the tendency to counteract and undermine the desired effect. And the net result? The only plausible answer here is that we cannot tell… and that is the point! It will depend on how the various forces work together and come into 'dynamic equilibrium', assuming that the system ultimately settles down to a stable pattern of behaviour. The aim of systems thinking is first and foremost to attempt to grasp the complexity of such organisational dynamics and the presence of possible unintended consequences. Drawing diagrams such as Figure 4 can be helpful, but it is the mode of thinking that is important.

There is currently much interest in system thinking of this kind in the UK public sector. Eileen Munro, for instance, has championed the systems approach in the context of child welfare (Munro, 2005). She led a national review of child protection services in 2010/11. Paraphrasing from the Review's interim report[2]:

> *A systems approach will help this Review to avoid looking at parts of the child protection system in isolation, and to analyse how the system functions as a whole. Social workers accept many previous reforms were well intended but their interaction and cumulative effect on frontline practice have had unintended consequences. The Review will use systems theory first to explain what has happened, providing a strong basis to build the Review's understanding. Second, the intention is to use systems theory to look forward, helping design an improved approach (p.10).*

'Systems dynamics' diagrams akin to Figure 4 feature prominently in the interim report, especially to highlight unintended adverse consequences ('ripple effects') of targets and excessive "procedural prescription" on professional discretion and staff morale. Munro's models illustrate how overemphasising compliance can reduce job satisfaction, leading to elevated staff turnover, which in turn will tend to reduce the experience level of staff, which itself will operate to lower the public status of the profession. And so unfolds the systemic

[2] The Munro Review of Child Protection Part One: A Systems Analysis. October 2010, Dept. for Education.

reading, showing how policy reforms aimed at enhancing practice, can produce opposing, counter-productive effects. This brings to mind another of McLuhan's gnostic "probes": "Every process pushed far enough tends to reverse or flip suddenly. *CHIASMUS* – the reversal of a process caused by increasing its speed, scope or size". Systems thinking draws attention to the potential for such paradoxical outcomes; it highlights the indeterminacy of human intervention in a complex world of diverse rationalities and incompatible goals.

References

Checkland, P. (1981). Systems Thinking, Systems Practice. Chichester: John Wiley.

Checkland, P. & Holwell, S. (2005). Information, Systems and Information Systems. Chichester: John Wiley.

Checkland, P. & Scholes, J. (1990). Soft Systems Methodology in Action. Chichester: John Wiley.

Flood, R. L. & Jackson, M. C. (1991). Creative Problem Solving: Total Systems Intervention. Chichester: John Wiley.

Munro, E. (2005). Improving practice: Child protection as a systems problem. Children and Youth Services Review, 27, 375-391.

Peters, T. (2005). Design – innovate, differntiate, communicate. New York: DK Publishing

Pfeffer, J. & Sutton, R. I. (2006). Evidence-based management. Harvard Busines Review, 63-74.

Senge, P. (1990). The Fifth Discipline: the art and practice of the learning organization. London: Century Business

Wastell, D.G. (2011). Managers as designers in the [ublic services: beyond technomagic. Triarchy Press.