Back to the future: digital decision making

Journal Item

Back to the future: digital decision making

Abstract

The process of making decisions about the conception, design, development, deployment and regulation of complex information and communications technologies (ICT) systems with the potential to effect significant changes in society could be labelled 'digital decision making' (or DDM for short).

DDM is not the rational process that we might assume or wish it to be. It can even be difficult to define the boundaries of the social, political or technical environments to which the process applies. It depends on craft knowledge, power and agenda, politics and situational messiness, personal values, law and environment, and a host of other factors starkly illustrated by cases ranging from the Three Mile Island to the space shuttle Challenger disasters. Too often DDM leads to information systems failures and it is time we started to learn from those past failures.

Introduction

This paper considers decision making related to and involving the development, deployment and regulation of complex information and communications technologies (ICTs). The shorter term ‘digital decision making’ or DDM (or ddm) will be used to represent this concept. Digital decision making is not a commonly used/recognised jargon phrase in the industry or academia. It should not be confused with some kind of artificial intelligence or the idea of machines making decisions. It is simply, in this context, a literary device to make life easier, hopefully, for the reader.

Information systems which have a considerable impact on public welfare are being created, deployed and regulated often without a fundamental understanding about what these systems are required to do, how they might actually work in practice or what their limitations might be. Electronic voting systems, child databases, DNA databases, national identity cards systems, national health information systems, air passenger profiling, motorist tracking systems, CCTV and communications data retention schemes are being deployed at a phenomenal rate. The whole area of information related laws is changing rapidly. All this is happening in piecemeal fashion without a genuinely coherent overall view or a sound understanding of what the technology can and, almost more importantly, cannot do.

Developments surrounding the Internet frequently affect things like privacy or access to education, for example, which are normally regarded as subjects for political concern. Yet debates about the development, regulation and deployment of the technologies are often regarded as technical or legal arguments about specialist subjects, rather than as matters which should concern everyone. The viewpoint implicit in this paper is that this is misguided.

Perhaps if DDM was considered like environmental decision making, it might get a higher public profile? If the last remaining trees and hedgerow in my neighbourhood were being ripped up to make way for a waste incinerator, I would feel strongly enough to speak to local people about it, sign a petition or write to my member of parliament. I would do it because I could see the impact it would have on my family's day-to-day quality of life. But we do not realise the effects of new technologies because they are not always so obvious or immediate. Neither do we know what powers we have to engage in decisions related to these new technologies.

This paper examines the way we go about making decisions about some of these complex technological systems in practice. It makes some recommendations about ordinary people and experts working together on complex DDM and concludes with an outline of a theoretical decision making methodology for the development, regulation and deployment of communications technologies.

How decisions are made

Everyone makes decisions from the moment they wake up in the morning until they go to bed at night. These range from the simple decisions to brush your teeth or take the usual route to work, to the more complex such as considering whether to change schools or jobs or what our responses to the threat of terrorism should be.

Know-how

In making decisions we go through a set of thinking and other processes consciously or subconsciously. When we scrub our teeth regularly it is to avoid the need for preventable dental treatment at some point in the future. We are prepared to invest some time and effort each day, along with the money we have spent on toothpaste and toothbrushes, to achieve this outcome. It is a rational decision but to some degree subconscious because it is just part of our daily routine.

Most of what we do is at the level of cleaning our teeth i.e. 'habit' or 'know-how' of one sort or another. Babies learn to walk, talk and recognise their mum and dad through an intense process of trial and error. Young children survive the shock of realising the universe does not revolve around them without reading a library of self-help books, when, for example, they first start having experiences outside of their immediate family, such as starting school. All the time they are making huge numbers of subconscious decisions in an attempt to get their muscles to move in a coordinated fashion or reacting to the feedback supplied via the complex social dynamics of the playground, in order to engage better the next time. They are constantly assimilating knowledge of the patterns of the world around them and developing skilled responses to enable them to survive and satisfy their curiosity and need for e.g. social interaction. Know-how in decision making gets automatically programmed into our brains through experience. It is valuable not just at the level of deciding to make a cup of coffee but also in much more complex decision making situations.

Think of the story of the old engineer called into to fix a ship's engine when the problem seems beyond company personnel. The engineer does a long and thorough inspection of the engine, takes out a small hammer and lightly taps it at one point, whereupon the engine magically springs back to life. He later sends the ship's owners a bill for £5100, which they complain about. All he had done was to tap the engine once with a hammer. So the engineer itemises the bill:

- Engine inspection £95
- Tapping with hammer £5

- Knowing where to tap £5000

The knowing where to tap is the know-how. It is difficult to define and even more difficult to measure it. Yet we know it when we see it – the skilled carpenter or sportsman, the successful business woman, the popular child in the playground. We do not often think about it as the ability just to get on with the multitude of life's decisions. We all use our own unique know-how all the time, consciously or not, in making the decisions large and small that get us through the day.

Rational decision making

There are a large number of different types of decisions and ways we go about making those decisions as individuals or as part of some group. In a rational decision making process we systematically follow a recognised series of steps to identify various options open to us and then choose one.

Benjamin Franklin [1753] once said that when he was faced with a difficult case, he would write down the pros and cons on a sheet of paper as they came to mind over three or four days. He would then weigh these against each other and decide on a course of action. That seems a rational thing to do.

Security specialist, Bruce Schneier [2003, p14-15] offers a five-step decision making process, which he believes applies universally to any decisions about security, including the regulation and deployment of technology for security purposes. He asks a series of questions:

1. What problem are you trying to solve?
2. What solution are you proposing?
3. How well does the solution deal with the problem?
4. What other problems does the solution create?
5. How much does the solution cost?

Project managers and engineers are familiar with another rational approach:

1. Survey the situation.
2. Specify the problem.
3. Identify a series of alternative options to tackle the problem.
4. Assess the alternatives.
5. Choose one and implement it.
6. Monitor the outcome and adjust action in accordance the relevant feedback.
7. If the 'solution' works move onto the next problem. If not start again.

These approaches represent variations on a theme which people use in professional or personal contexts. They all rely on the gathering and assessment of accurate information or facts about a situation in order to make a rational choice about the best course of action. Because the decision situation is usually in a state of flux we often find we have to go through the steps more than once. So if a parent decides to get a child a computer games console for Christmas and a more modern version becomes available, then the decision situation has changed and the decision process needs to be re-visited.

Complexity in decision making: garbage can situations

These rational approaches are much richer than a superficial list of the steps involved will make them appear but they have been criticised as ignoring or underestimating the complexity and real-world uncertainty and confusion involved in actual decision making.

James March[1] [1994, p198-206] has written that real-life decision situations are often better characterised by the 'garbage can' metaphor than artificially rational steps. Many decision environments display fuzziness – complex interactions between the people, technologies, problems, solutions, social, organisational, legal and economic contexts. These are all mixed up together in a garbage can at a point in time and the relationship between a problem, a solution and a decision maker may have more to do with them coming together in the same place at the same time, than any rational process.

March says that people constantly have a range of issues, professional and personal, competing for their attention. The deadline for contract negotiations on a big project is looming; your partner is suffering from a de-generative illness; one of your kids is being bullied at school; there is a public transport strike so you were late for work; there's a major last-minute hitch with the contract but a colleague thinks a new computer system will fix that; some members of your team are unreliable; there's a fly buzzing around the conference room. In complex decision environments it is impossible to know or shut out everything but the relevant issues, then analyse these through some rational process to come to the 'right' decision. Real life is much messier than that.

Bounded rationality in decision making: satisficing

So if the real world is so messy and there is too much extraneous noise in complex decision making situations to act entirely rationally, what can be done? Well we could apply a rational approach to the limited amount of apparently relevant data we can extract from the situation. If the assessment of a variety of bidders for a government telecommunications contract suggests that two of the companies could meet the requirements within the required budget, then randomly picking one through the flip of a coin would lead to a 'good enough' choice.

When I got my latest mobile phone, told the shop assistant what I was looking for, was shown two matching that specification and randomly picked one. It does the job I need it to do, most of the time. Occasionally the battery runs out at inconvenient times. Decision theorist, Herbert Simon [1957], coined the term 'satisficing' for this partly rational, just-good-enough approach to decision making.

The British radar technology in World War II was inferior to that of the Germans, so much so that when the Germans captured a British radar set in 1940 it was declared so obsolete as to be useless.[2] The technology, however, was good enough, as part of an integrated system, to collect the raw data on approaching enemy aircraft. This raw data from their chain of radar stations and visuals from the Observer Corps was passed on (via radio telephone and teleprinters) to headquarters and an integrated set of operations centres, where it was assessed, filtered, analysed and turned into useful information at varying levels. This then facilitated the scrambling of the right fighter squadrons and even more specific instructions to be radioed to the RAF pilots once in the air, to enable them to intercept their enemy at the earliest opportunity.

The Germans had better information technology (radar). The British had the better information system (radar, human intelligence, signals intelligence, and an integrated, purpose-developed system, allowing the situation to be viewed holistically, as well as delivering the right information to the right users, at the right levels, in a useful format and in sufficient time to act on it). The better information system prevailed in the Battle of Britain in 1940 and it got built in time largely due to the decision of those involved to use technology that was just good enough to get the job done.

Factors that influence [digital] decision making

Amongst the factors that influence decision making, the personal values and relative power basis of key decision makers are fundamental.

Facts, values and power

It is important to realise that we are all conditioned and predisposed to believe certain stories more than others. This is because of our individual prejudices and values. If someone doesn't like George Bush or Gordon Brown, a song that makes fun of them will appeal. If we do like Bush and Brown, however, we might find the song offensive.
The Royal Commission on Environmental Pollution [1998, Chapter 7], defined values as follows:

> "We understand values to be beliefs, either individual or social, about what is important in life, and thus about the ends or objectives which should govern and shape public policies. Once formed such beliefs may be durable."

So values are strongly held personal beliefs about what is important and about how the world ought to be.

Arguments about DDM can be very heated, as the public debate about ID cards and human rights has shown. Those in favour and those against ID cards both argue passionately from the perspective that they are 'right' and the other side is 'wrong'. Both the pro- and anti-ID card factions cite facts and statistics in their defence. To get to the bottom of some of the confusion, it is helpful to distinguish between facts and values.

A fact is something that we believe to be objectively true. A piece of gold weighs more than a similar-sized piece of aluminium. Water boils at 100 degrees centigrade (at standard atmospheric pressure). The price of *The Times* newspaper in the UK was 70p on 17 June 2008. These are all things that can be observed and measured by processes that are not subjective. They can be agreed upon by all reasonable people.

A value, on the other hand, is a belief that something is good or bad. That the music of Cole Porter is better than that of Mozart; that ID cards are good; that euthanasia is always morally wrong.

It could be said that facts are beliefs about what is, and values are beliefs about what ought to be.

How is this distinction helpful? Well, because many public arguments involve a mix of facts and values. They are presented as if they are disputes only about facts, whereas they are really about conflicts in values. This is significant because disputes that are about facts can, in principle, be resolved by some objective process that can establish which assertions are factually correct. You can imagine a kind of impartial court that could adjudicate between the rival claims and reach a judgement acceptable to all.

But conflicts about values cannot be resolved in this way. There is no purely objective process by which the dispute can be resolved. There is no rational process by which someone who believes in euthanasia can convince someone who is opposed to it.
So when we examine DDM situations we should try to distinguish between their factual content and their value-laden contexts. The balance may determine whether or not they are resolvable by some rational decision making process. Secondly, we should remember that the disputes between the protagonists in DDM situations often arise out of differences in values.

Winston Churchill abolished identity cards in 1952. In the wake of the 11th of September 2001 attacks on the US, the then Home Secretary,[3] David Blunkett, embraced a plan to reintroduce them. He invested a lot of energy in pursuing it, as have his successors, Charles Clarke, John Reid and Jacqui Smith. The proposal for the high tech system came about at a time (2001/2) when the government was facing serious questions on terrorism and immigration. The idea appealed to Mr Blunkett, someone with a strong belief in the need for government to be taking big decisions to tackle big problems. He was also in a position, at the time, to make it happen. This plus the fact that terrorism and immigration are incredibly complex issues fits the temporal link theory in March's garbage can process.

If a government minister, or anyone else, strongly believes some action is the right thing to do, it is difficult to get that person to question that belief. The most powerful actors also tend to have the means to act on their beliefs.

There is an extra complication in the context of powerful actors. People like prime ministers, presidents and chief executives tend to be surrounded by people whose jobs depend on keeping the boss happy. They therefore have an incentive to tell the boss what she wants to hear i.e. to reinforce her beliefs. Churchill put it like this:

"The temptation to tell a Chief in a great position the things he most likes to hear is the commonest explanation of mistaken policy. Thus the outlook of the leader on whose decisions fateful events depend is usually far more sanguine than the brutal facts admit."[Jones, 1998, p161]

The good leaders know this and compensate accordingly.

It is important, however, to understand the power dynamics, the personal values and the agendas of the most powerful actors in any DDM situation.

Thinking traps

The thinking trap can be a barrier to even bounded rationality in decision making. Geoffrey Vickers [1970, p15] described it thus:

"Lobster pots are designed to catch lobsters. A man entering a man-sized lobster pot would become suspicious of the narrowing tunnel, he would shrink from the drop at the end; and if he fell in he would recognise the entrance as a possible exit and climb out again – even if he were the shape of a lobster. A trap is a trap only for creatures which cannot solve the problems that it sets. Man-traps are dangerous only in relation to the limitations of what men can see and value and do. The nature of the trap is a function of the nature of the trapped... we the trapped tend to take our own state of mind for granted – which is partly why we are trapped."

He notes that we can only start to climb out of our self-made thinking traps when we recognise that we are *in* a trap and start questioning our own limitations and the assumptions that led us there.

I used to find it hard to accept that lawyers were prepared to act for people or organisations who had allegedly engaged in reprehensible practices. Yet it is a fundamental tenet of a just society that people accused of even the most heinous crimes are entitled to a fair trial.[4] This thinking frame – 'only good people should be entitled to legal representation' – represented a trap in my thinking inhibiting a wider understanding of the legal system.

Proxies in decision making

Because it is impossible for us to do everything or understand every complex situation we face, we often employ proxies to make decisions for us. A proxy is a person, an organisation or a machine that acts on our behalf in some way.

We vote for politicians who subsequently sit in parliament where our laws are passed. The English Football Association appoints the England manager to pick the team to play in the World Cup. Organisations have proxy computers that act as gatekeepers between the company network and external networks connected to the Internet. The chef at the restaurant sources the ingredients in the food customers are served. Law enforcement authorities and intelligence services are our proxies in fighting serious crime.

Proxy decision makers present us with a problem, however. Even though they are making decisions on our behalf, we may or may not trust them. Governments are often reported as being untrustworthy in the eyes of the public, especially in the wake of political scandals. Proxies have to earn our trust through success, transparent decision making, third party audits, experience, know-how and recommendations of people we do trust. Trust in governments is dependant on transparency.

Proxies will not necessarily make the decisions we ourselves would have made faced with the same circumstances, since they have their own complex agendas, motivations and constraints.

Complexity – the technology: Three Mile Island

At the heart of computer technologies lie hardware with millions of tiny electronic components and software programs with millions of lines of code, which together constitute some of the most complex machines that have ever been built. That very complexity is a key factor in the success or failure of DDM processes involving these machines.

Charles Perrow[5] [1999, p304] thinks that some complex technologies and the complex systems of which they form a part, such as nuclear power plants, are so prone to failure with catastrophic effect that we should abandon them completely. Perrow describes the partial meltdown of the reactor core at the Three Mile Island nuclear power plant in 1979 as a 'normal accident', the inevitable result of the complexity of the plant system, and the tight coupling of its component parts.

The complexity means no one can fully understand the system and the tight coupling means that failure in one component can have a ripple effect, leading to a string of other components failing like dominoes falling over. The complexity also leads to parts of the system, including the human actors,[6] interacting in unexpected ways (because they are interlinked in unexpected ways) resulting in the emergence of properties of the system which would not have been predicted in advance.

At Three Mile Island part of the cooling system had been isolated for some maintenance. It was standard practice to use compressed air to clear a blockage. The blockage proved to be stubborn and difficult to shift and the operation was taking much longer than usual. A small amount of water leaked back through the compressed air pipes into the control instruments triggering a shut down of one of the plant's three main cooling systems and of the electricity generating turbines. A stuck pressure relief valve in the reactor core cooling system then went undetected partly because of misleading and hidden indicators in the plant's control room.[7] Operators in the control room were left with the erroneous impression that pressure was building up dangerously in the reactor core cooling system, which if it failed would leave them with no means of cooling the reactor and preventing a total meltdown. So instead of pumping more cooling water into the system they drained water away, in order, so they thought to prevent the core cooling system failing catastrophically. It is difficult to imagine the stress endured by plant operators faced with a nuclear disaster and a power plant system behaving in ways they could not understand despite their significant combined level of experience.[8]

It was not until more than two hours later, when a new shift supervisor, Brian Mehler, arrived on the scene, that the problem with the valve was discovered and they began to pump more water into the system to prevent a disaster. Mehler modestly says he merely "brought a fresh pair of eyes to the room" but he was able to enter a highly stressed environment and test his theory about the valve to a natural conclusion. His colleagues had also considered the valve as a potential problem but within a couple of minutes of the start of the incident over one hundred alarms were going off in the control room. In the confusion of frenzied activity, a temperature reading on the valve had been either considered to be within the required limits or reported erroneously to the people in charge.

This again was partly down to serendipity. The pressure valve was known to have a small leak which could not be easily fixed, so the computer linked to the temperature indicator on the valve line had been programmed not to give any readings over a specific limit, 280°F. Mehler noted the temperature, still felt it was unnecessarily high and asked for the valve to be isolated. Almost instantaneously the system began behaving in predictable fashion and they were able to bring the water levels up thereby avoiding a disaster.[9] According to a US Presidential Commission [1979] report on the accident the nuclear core had been less than an hour from total meltdown.

Complexity – the situation: Challenger

I have drawn attention above to some of the key factors influencing decisions – personal values, relative power, thinking traps and the complexity of the technology but there are quite a number of others which I would group together under the heading 'complexity of the situation'. These include:

- The decision makers
- Decision criteria
- Time
- Dynamic (changing) nature of the situation
- People affected
- Law
- Decision making models (such as cost-benefit analysis etc)
- Decision environment (organisational, ecological, economic, social, political and physical).

Consider the Challenger space shuttle disaster at the Kennedy Space Center at Cape Canaveral in Florida, on 28 January 1986. The technical cause of the accident was the failure of rubber O-ring seals in one of the booster rockets. The freezing temperatures at the launch meant that the rubber was not capable of doing the sealing job required. Escaping gas destroyed one of the key fixtures securing the booster rocket to the main fuel tank and burned a hole in the side of the tank. The out-of-control rocket swivelled around its upper fixture, crashing into the top of the fuel tank and leading to a massive fireball. The space craft broke up. It was just 73 seconds into the flight. Engineers at Morton Thiokol, the company which made the booster rockets had strongly advised against launching in those temperatures and company managers, as well as those at NASA, were later vilified for acting against this advice.

The launch had nearly happened the day before the accident, when technical problems led to it being abandoned during countdown and reset for the following day. Shortly thereafter, at NASA's request, Morton Thiokol engineers had a meeting about possible problems with the performance of the O-ring seals in the freezing temperatures forecast for the next day. There was a history of hot booster gases burning through O-rings, the most significant damage occurring on a shuttle flight in 1985, when the launch temperature had been the lowest on record, 53°F. Morton Thiokol engineers and management agreed they should not sanction a flight below this temperature. At a teleconference later that evening, however, under pressure from NASA to agree to the launch, Morton Thiokol took a 'management decision' to agree it should go ahead, in the face of their engineers' objections.

John Young, NASA's chief astronaut, in an internal memo following the accident said:

"There is only one driving reason why such a potentially dangerous system would ever be allowed to fly – launch schedule pressure."

NASA was regularly criticised and ridiculed in the media and by politicians for launch delays and excessive spending. This particular flight had drawn a lot of media attention from all over the world because it was to include the first teacher in space, Christa McAuliffe. It does seem unlikely, though, with the attention of the world's media more intense than it had been for many years that NASA managers would have risked the flight, if they had any serious doubts about its safety.

Diane Vaughan [1997] characterises this misplaced confidence in the safety of the mission, in spite of the clear technical advice to the contrary, as a 'normalization of deviance'. She tells a convincing story of how the history of NASA has been one of budgetary constraints which led to design trade-offs in the shuttle they would have preferred to avoid. In spite of the fateful decision, which with hindsight proved to be so disastrous, she also discovered many examples of cases where NASA managers had made very expensive decisions purely in the interest of safety. Crew training, launches frequently abandoned on safety grounds in spite of launch schedule pressure, huge numbers of complex procedures and safety checklists and the fact that they talked at length to Morton Thiokol on the eve of the launch point towards an organisational culture which clearly did not neglect safety.

Critically, after previous problems with the O-rings, the booster rockets had been tagged with a formal NASA 'launch constraint'. This meant the O-rings were a recognised safety concern serious enough to prevent a launch. Critically also, NASA had developed a formal 'waiver' procedure – a procedure that allowed NASA personnel to ignore normal rules and procedures, when they needed to. Under the waiver procedure five shuttle missions had proceeded, even though the problems with the O-rings were known.[10]

In these circumstances it is possible to see a false confidence in the safety of the O-rings developing. The argument is that it has not failed catastrophically in the past, so it will not do so the next time either. Hence Vaughan's conclusion that NASA slowly evolved into a state where they had actual formal procedures allowing crucial safety issues to be ignored. This she characterised as the normalisation of deviance. That any organisation should draw up procedures to bypass other formal organisational procedures, particularly those involving safety, might seem strange but it is extremely common. It is a well known, ironically unwritten rule of every organisation that the way to bring the place to a grinding halt is to work to the letter of organisational procedures.[11] This is why 'work to rule' is one of the standard tactics in the armoury of any union involved in an industrial relations dispute.

The Challenger shuttle type of situation always has multiple causes beyond the immediate technical failure or series of failures (in this case the O-ring, rocket fixture, out of control rocket, disintegration of shuttle). The organisation rationalised, and then tolerated serious safety problems due to launch schedule pressure, arising from the prevailing social, organisational, political and economic environment. The disaster points to the immeasurable importance of informed decision making at the heart of complex systems.

Social technologies

NASA's procedures allowing a shuttle launch to go ahead in spite of clear safety concerns could be considered to be a subset of 'social technologies'. Social technologies involve people, organisations and practices and mental and administrative frameworks and models for understanding situations, including language and numbers. They are often invisible and followed without question or awareness of their origins, or the need for contextual understanding because they form the fabric of our daily routines. NASA bypassed their safety procedures because it was routine, so routine in fact that they had established a formal process for doing it.

Social technologies include laws, organisational procedures and rules to regulate behaviour. They may also include the architecture of the technologies we are dealing with [Lessig, 1999]. In the case of digital technologies this may for example include built-in digital locks or digital rights management (drm) devices preventing them from being used in certain ways. Social technologies can structure how we think and act and therefore determine how decisions are made. A widely used (and abused) mathematical social technology is cost benefit analysis.

In the context of language, control of the language used in a decision process can be the key to controlling the outcome of that process. Language is rarely neutral in complex decision making situations. 'Intellectual property', for example, is considered by some to be something of a misnomer, which might be more accurately described as 'temporary and limited intellectual monopoly'. Describing someone as a 'citizen' or a 'consumer' subtly defines their role. Social technologies therefore include the mental structures through which we view the world and hence we come full circle again to the personal values that shape our thinking.

Informing DDM: Systems thinking

'System' is a word which is widely used to mean different things. When using the word system we are usually referring to something complex, with many interconnected components or subsystems that [should] work together in some coherent way. An information system, an education system, a legal system, and a health system are common examples. 'Systems thinking' is a way of letting us look at something as a whole. The whole system can be greater than, less than or equal to the sum of its parts.

The England football team does less well than would be expected given the proven abilities of the individual players at its disposal. A group of managers, with individual IQs above 120, were found to have a collective IQ below 63, when observed working together on a particular task [Senge, 1992]. The Challenger space shuttle had over a thousand subsystems or components with 'criticality 1 waivers', the failure of any one of which would have been enough to lead to the loss of the shuttle and the death of the crew. The UK government are introducing an identity system not just an ID card.

The idea is that by looking at the whole system we can better understand the complexity and interconnections between the system components and the system's overall effectiveness.

The very act of thinking of something in a different way is often enough in itself to enlighten our perspective of a particular system. A copyright system designed to maximise the income of copyright holders might require copyright to last forever. A copyright system designed to promote progress in science and the useful arts, however, might find never-ending copyright interfering with the ability to use other people's work to develop new ideas.

One of my favourite stories illustrating the power of viewing the whole system (in the context of its environment) and understanding multiple perspectives on a situation is the parable of the blind men and the elephant [Saxe, 1873].

It was six men of Hindustan
To learning much inclined,
Who went to see the Elephant
(Though all of them were blind)
That each by observation
Might satisfy the mind.

The first approached the Elephant
And happening to fall
Against his broad and sturdy side
At once began to bawl:
"Bless me, it seems the Elephant
Is very like a wall."

The second, feeling of his tusk,
Cried, "Ho! What have we here
So very round and smooth and sharp?
To me 'tis mighty clear
This wonder of an Elephant
Is very like a spear."

A third thinks it is a snake having touched a wriggling trunk, a fourth goes on to say it is like a tree, after touching the knee, a fifth touching the ear says it is like a fan, and a sixth says it is more like a rope after touching the tail. All six go onto argue "loud and long" and

Though each was partly in the right
and all were in the wrong.

So oft in theologic wars,
The disputants, I ween,
Rail on in utter ignorance
Of what each other mean,
And prate about an Elephant
Not one of them has seen!

Systems' thinking was pioneered by biologists who stressed the need to consider living species as integrated wholes. Life is greater than the sum of its component parts. James Boyle [2003] is essentially urging the duck hunter, the bird watcher and the family whose water supply is contaminated with industrial waste to use systems thinking or to think 'systemically', when pointing out how the term 'environment' enables them to see their shared interests. They might not like each other – especially the duck hunter and the bird watcher – but they have a shared interest in protecting the birds' environment. Thinking about the environment at different levels facilitates understanding. Getting too deeply immersed in our own particular concerns in our own part of a system can blind us to the wider picture, which itself can suggest connections and help we would not otherwise see.

A system is dynamic i.e. it changes over time. It also has a purpose. This notion of purpose is incredibly important when dealing with information systems. Many information systems fail for the simple reason that the purpose of the system is never really clarified. Even if we do not know the purpose of a system it is useful to examine it as if it had a particular purpose. The intellectual property system may or may not have become something that favours particular vested interests, as some argue[Lessig, 2001; Boyle, 1997]. But it is still helpful to look at it *as if* it had the purpose of increasing access to knowledge.

The fundamental value of systems thinking and practice is that it facilitates the development of *different perspectives*. Simply thinking about a situation in a different way can be enlightening.

Messes and difficulties

One of the gurus of systems thinking, Russell Ackoff [1998], gave the term 'mess' a particular meaning in the context of decision making:

> "What decision makers deal with, I maintain, are messes not problems. This is hardly illuminating, however, unless I make more explicit what I mean by a "mess". A mess is a set of external conditions that produces dissatisfaction. It can be conceptualized as a system of problems in the same sense as a physical body can be conceptualized as a system of atoms."

Decision making situations are often unstructured. This contrasts with simple 'difficulties' where it is easy to identify and solve a particular problem.

Messes tend to be big, complicated, involve more people, more organisations, happen over a longer time scale, have serious impacts, confusing features, no obvious solutions and no real clarity often about what the problems are.

By contrast, difficulties are easier to pin down. The problem is clear, as are various solutions. When the car needs a service, I can cycle to work. The problem is bounded, small scale and clear.

Most of the DDM situations dealt with by government and regulatory bodies are of the mess variety. Yet they are often treated by policymakers as being difficulties. There is a misplaced faith in the magical ability of technology to solve problems in a way which turns a mess into a mere difficulty. This belief is particularly widespread amongst decision makers who do not understand the technology. Yet technology is just a tool. It will not miraculously evolve towards a state where it will comprehensively address an ill-defined mess. Yet policy makers rarely seem to understand this.

That is partly why we now have laws which protect the digital fences (or digital rights management, DRM) behind which digitised copyrighted works are locked. There is no law which makes it a crime to jump over a real fence. There is a law of trespass which but the fence is not considered to need its own special legal protection in order to deter someone from jumping over it. Likewise in the digital copyright context there is a law of copyright to protect the work behind the digital fence. Yet the digital fences also get their own special laws and the penalties for breaching the fence are much more severe that those for breaching the copyright the fence is supposed to protect; as the arrest and jailing of Dimitry Sklyarov showed.[12] The state of the intellectual property system is a mess of the Ackoff variety. It is a mess partly because it is now applying what historically was industrial regulation, at the level of the individual.

Informing DDM: Experts and ordinary people

I believe experts and ordinary people have a lot to learn from each other in dealing with DDM situations, especially where the DDM process has been captured, as in the case of the intellectual property system, by a relatively small number of interested parties [Drahos and Braithwaite, 2002]. Consider a story where the experts were wrong.

Shirley McKie and the fingerprints experts

In January 1997, Shirley McKie was one of a team of police officers investigating the murder of Marion Ross in Kilmarnock. David Asbury was soon arrested as a suspect in the murder based on fingerprint evidence found at the scene. Today fingerprints are widely accepted as an infallible means of personal identification. The fingerprint identification process can go wrong though and it did in the Ross murder case.

In addition to the fingerprints at the scene implicating David Asbury in the murder, another print was found there, which four fingerprint experts at the Scottish Criminal Records Office identified as belonging to officer McKie. Since McKie had never been in the house she suggested there must have been a mistake. The main case against the murder suspect, Asbury, amounted to the fingerprint evidence, however, so senior officers believed that it must have been McKie who had made the mistake. Not just one but four experts had confirmed the identification of the print. Whereas it might have been possible for one examination to be in error, it was highly unlikely that four experts using this long-established scientifically reliable identification technique could be wrong. Or was it?

Officer McKie then came under pressure to change her story since all the fingerprint evidence would come into question if the experts were seen to have made a mistake with the print identified as hers. McKie refused to yield to the pressure because she knew that the print could not have been hers. Unbeknownst to McKie at this stage four other fingerprint experts at the Scottish Criminal Records Office had also examined the disputed fingerprint and had refused to confirm it was hers.

During the murder trial of David Asbury, McKie denied under oath on the witness stand that the fingerprint belonged to her. Asbury was convicted of murder and McKie was subsequently prosecuted for perjury. In her defence she employed two overseas fingerprint experts who confirmed the disputed print could not belong to McKie and explained in some detail why this was the case. Their clear explanation, accessible to ordinary people, has since been shown in several BBC TV programmes on the case. In May 1999, McKie was unanimously acquitted of perjury and the judge, unusually, commended her for

"…the obvious courage and dignity which you have shown throughout this nightmare… I very much hope you can put it behind you. I wish you all the best."

Three years later David Asbury's murder conviction was quashed by the Appeal Court, which agreed that the fingerprint evidence against him was unreliable.
In 2000, after a lot of campaigning and a couple of BBC programmes, an investigation by Her Majesty's Inspectorate of Constabulary concluded that the fingerprint mark could not have been made by McKie. Hundreds of fingerprint experts from all over the world have now examined the evidence and come to the same conclusion. There have been numerous investigations and reports, TV programmes and in 2006 a parliamentary inquiry in the Scottish parliament.

The process surrounding the fingerprint science in the Scottish Criminal Records Office went seriously wrong in the McKie case and the investigation of the Ross murder.[13]

By the autumn of 2006 the four fingerprint experts at the centre of the case were coming under pressure to resign or retire but that was scapegoating. There was a complete failure of the social technologies – the processes, procedures and management – in the Scottish Criminal Records Office, the police and the Crown Office and Procurator Fiscal Service (COPFS) which brought the McKie prosecution. The system surrounding the science and the scientific experts failed and failed badly, when those experts made a mistake.

The McKie case is a clear indication that the application of scientific expertise to decision making processes should never be accepted with blind faith, especially if there is clear evidence of a professional difference of opinion between the experts, which subsequently gets hidden from public view. The reliability of the expert opinion crucially depends not just on the science but the integrity of the organisational, social, legal and technical systems supporting and surrounding it.

Experts can and do make mistakes, all the time. They are only human after all and, hard as it is for most people, in some ways it can be even more difficult for an expert to admit a mistake. The reputation of the expert and the system are vested in their opinion and the admission of a mistake might well put a dent in the reputation of both. In addition it can put a dent in the public's confidence in that whole field of expertise and any associated systems of which it forms a part.

In the Shirley McKie case, the experts' mistakes were backed up by the full weight of the Scottish criminal justice system. Even when clear evidence of the existence of the mistaken expert judgments was provided, the entire system found it nigh on impossible to admit the errors.

Yet an expert's mistake does not necessarily mean the science is wrong. It just means that the expert made a mistake. The general public should, therefore, be alert to the possibility of such mistakes, through critical questioning of expert opinions and understanding of the mistakes when they happen, to help ensure they can be accepted and corrected. The systems which require those expert opinions as part of their ongoing operations should be equally ready to understand and rectify such mistakes when they occur. To the extent that such mistakes expose serious systemic, management or structural problems in a system, the people involved should be prepared to take the necessary steps to ensure those systems and structures are put right accordingly.

Experts and Economic models for DDM – cost benefit analysis

Most big organisational or government decision making situations are dominated by one factor – money. How much is it going to cost and how much is it going to earn or save us? We calculate the costs and the benefits of the various options and choose the one with the highest net benefits.

The modelling of this is derived partly from the simple and compound interest calculations we all should have learned in primary school. £500 invested for two years at 10% simple interest per annum, earns £50 in interest after the first year and a further £50 after the second year, or £100 in total. Most financial institutions use a compound interest formula on borrowings or investments. So the interest earned each year is added to the original investment, for example and interest for the following year is based on this accumulated sum. So in the case of the example above, the interest would be £50 after the first year. The second year starts with a total investment of £550, earning further interest at the end of the year of £55. So the total interest earned under the compound interest system is £105. The investor is £5 better off than if the bank were using a simple interest formula.

£100 today is worth more than the same £100 a year from now because of the interest that could be earned from investing it in the meantime. On big projects money flows happen over a period of years. One way to get an assessment of these costs and benefits is to reduce all future flows to their present values, using discounted cash flow analysis. This is an inverted compound interest calculation. Standard, usually computerised, cost benefit analysis calculations for big projects have to use estimates of how much money is coming in and out at various times. In addition they need assumptions of the likely prevailing interest rates over the course of the development and operation of the system.

The kinds of assumptions that underlie all these estimates can be of epic proportions, which is why we should not be surprised to read regular media condemnations of cost overruns of major government information systems projects. We should actually be more appreciative of the vast numbers of big projects [though there are not many of these in the IT sector] that run to budget and planned timescales, but which do not come to our attention because people doing a good job is not a newsworthy story.

The benefit of the computer model is that it can quickly provide estimated outcomes for best-case and worst-case scenario estimates to help with our decision making. Use of cost benefit analysis computer models is an almost compulsory part of the planning process for government projects. But we always need to be aware that the output of the computer models is only as sound as the data fed into them and the validity of the assumptions underlying the construction of the model.

Discounted cash flow modelling can give us a surprisingly powerful insight into the value of intellectual property. Copyright typically lasts for the life of the author plus 70 years. So if a copyright lasted 110 years, assuming a discount rate of 10% per annum and a constant income, the copyright holder earns 99.997% of what she would earn if the copyright term lasted forever.[14] So we need to be sceptical of arguments declaring that the only way intellectual property owners are going to derive an appropriate income from their creative assets is by further extending the term [Kretschmer et al, 2008].

The EU's Internal Market Commissioner, Charlie McCreevy, is committed to extending the copyright term in sound recordings from 50 to 95 years.[15] In real terms, this would make a difference in the income over the life of the copyright of a fraction of a percentage point. In the case of recordings made in the 1950s on which the copyright is about to expire, however, it would instantly provide an extra 45 year monopoly to companies or individuals holding those rights.

So discounted cash flow models are useful. Likewise other forms of modelling can be very powerful tools when built and used appropriately and with integrity. But they are just tools. Ultimately the decision making is done by people and those people need to understand the limitations of their tools.

Ordinary people and models: critical questions

Ordinary people have a crucial role to play in questioning the simplifications and assumptions that go into all these decision making models. Everyone in DDM situations has their own perspective and agenda. We need, therefore, to be careful about accepting arguments at face value.

When analysing a DDM situation, try to consider the situation from each stakeholder's perspective and understand their agendas. Identify explicitly any underlying assumptions that they may be making and assess their credibility and that of the arguments derived from them. Aviation pioneer, Howard Hughes, made some of the most important commercial decisions of his life by instinct or gut reaction [Higham, 2004] so this might not always be possible but it is important to try. Think about the system's purpose and whether any of the stakeholders are focused on it. Analyse and evaluate the various claims, information and research. Even if we do not feel competent to evaluate research or models beyond our field of competence we can still ask the simple questions:

- What is the purpose of the system? (It is ridiculously common for there to be no clear consensus on this, even amongst the key stakeholders).
- Who paid for the research?
- What are the underlying assumptions of any models used?
- What are our alternative courses of action?
- Are there other alternatives that have not been considered?
- How well does our proposed action address the purpose of the system and from whose perspective?
- What other problems is the system likely to create?
- What is the likelihood of these?
- Are the new problems/messes better or worse than those we have set out to address?

The DDM system will exhibit emergent positive and negative consequences. Can we anticipate any of these and exploit or protect ourselves from them? How can we monitor and improve the system iteratively, in response to what we learn?

How might the system work in a 'best possible scenario' case and a 'worst possible scenario' case? Round off with Schneier's [2003] final two questions:

- How much does it cost?
- Is it worth it?

Do not be distracted by the 'it's what the public wants' argument. Just as in science, where truth is determined through the scientific method and rational and ethical peer review, the prevailing state of a reported majority of public opinion should not necessarily be the determining factor. Informed and critically aware public opinion does, however, have a key role to play in DDM.

A DDM framework

To round off I offer a theoretical digital decision making (DDM) framework. It is based on some Open University work I have been involved in, producing and delivering masters degree courses in the area of environmental decision making[16] and my contention that ordinary people need to be involved in the DDM process.

The frameworks developed in the Open University make the assumption that many environmental decision making situations are unstructured messes – complex systems of problems all mixed up together – rather than easily identifiable simple difficulties. Before we can outline solutions to perceived problems, we need to make an effort to understand these messes and the multiple perspectives different people have on them. Through understanding we can begin to identify the real problems in the situation and then start thinking about how to deal with them. Decision making is also seen as an ongoing learning process involving interested or affected parties, often called 'stakeholders'.

DDM situations all too often begin with imagined technological solutions to complex messes instead of attempts to understand those messes. ID cards and passenger database computer profiling systems are the solution to terrorism; the children's index database[17] the solution to the welfare of vulnerable children; changes to intellectual property laws the solution to the mass copying capacity of the Internet. If we ask how well these solutions address the problems they are apparently targeted at, the answer is, inevitably, not very well if at all.

My proposed DDM framework is shown in Figure 1. It starts with the need to understand and to some degree provides a 'how to' guide for decision makers. The framework, however, is a modest proposal and should not be considered to be the one-true-guide to DDM.

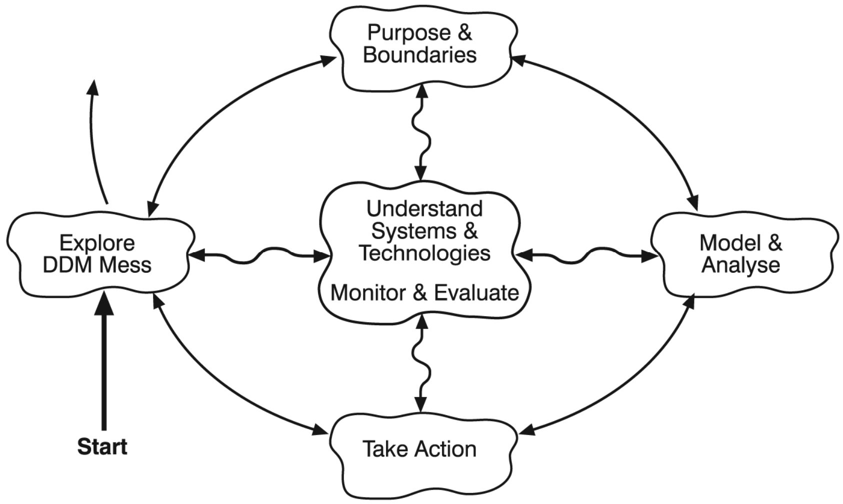

Fig. 1 Digital decision making framework

The idea behind the framework is that it provides a structure to think about DDM. The cyclic shape is a deliberate attempt to generate the picture of DDM as an ongoing learning process. Although there are sometimes clear and specific decisions to be made at the 'take action' stage, this is not the end of the process. Frequently, even when information technology projects are well planned, the whole thing falls apart in the implementation of the system, because the decision to go ahead with the project is the cut-off point. Actually building or deploying the system becomes someone else's problem and the demanding reality of deployment and operation is not taken into account in the resources set aside for implementation.

Hence the suggested framework includes the decisions, at the 'take action' stage, as one element in a continuous learning process. The decisions are crucial but the cycle does not suddenly stop at that point if the systems are to be successfully deployed or regulated. Take a look at the various stages of the suggested DDM framework:

- Explore the DDM mess
- Be clear about your purpose and your boundaries
- Use models, where appropriate, to help analyse, understand and frame the situation and the feasible and desirable options
- Take action
- Understand the systems and technologies; and continually monitor and evaluate proceedings at each stage of the process.

Not surprisingly the need to understand the systems and the technologies is at the centre of the framework, since I have suggested that this is the key to improving DDM processes. Every other stage of the framework is linked to that centre by double-headed arrows indicating a two-way flow of information.

The various activities will not necessarily always happen in the neat sequence that a two-dimensional framework is restricted to portraying. The very nature of a messy DDM situation means it is rarely neat or sequential, though it may be possible to identify emergent patterns in the mess.

The various stages of the framework are represented by blobs rather than neat, sharp shapes like ellipses or rectangles, in order to emphasise the dynamic nature of the process. The stages are not simply neat, fixed and sequential. Each can change shape like an amoebic micro-organism and be active in parallel with all the others, as the developing of under-standing and the operational learning process proceeds.

The clearest visual distinction between this framework and other decision making methodologies is its continuous cyclic or iterative nature, emphasising the importance of the learning process. Other approaches have distinctive starting and end points, though Schneier [2003] too stresses the importance of appreciating complexity and thinking about security decisions as part of a continuous learning process. Some people find it helpful to think of this cyclic continuous process as a cascading series of cycles round the frame-work, as shown in Figure 2 below.

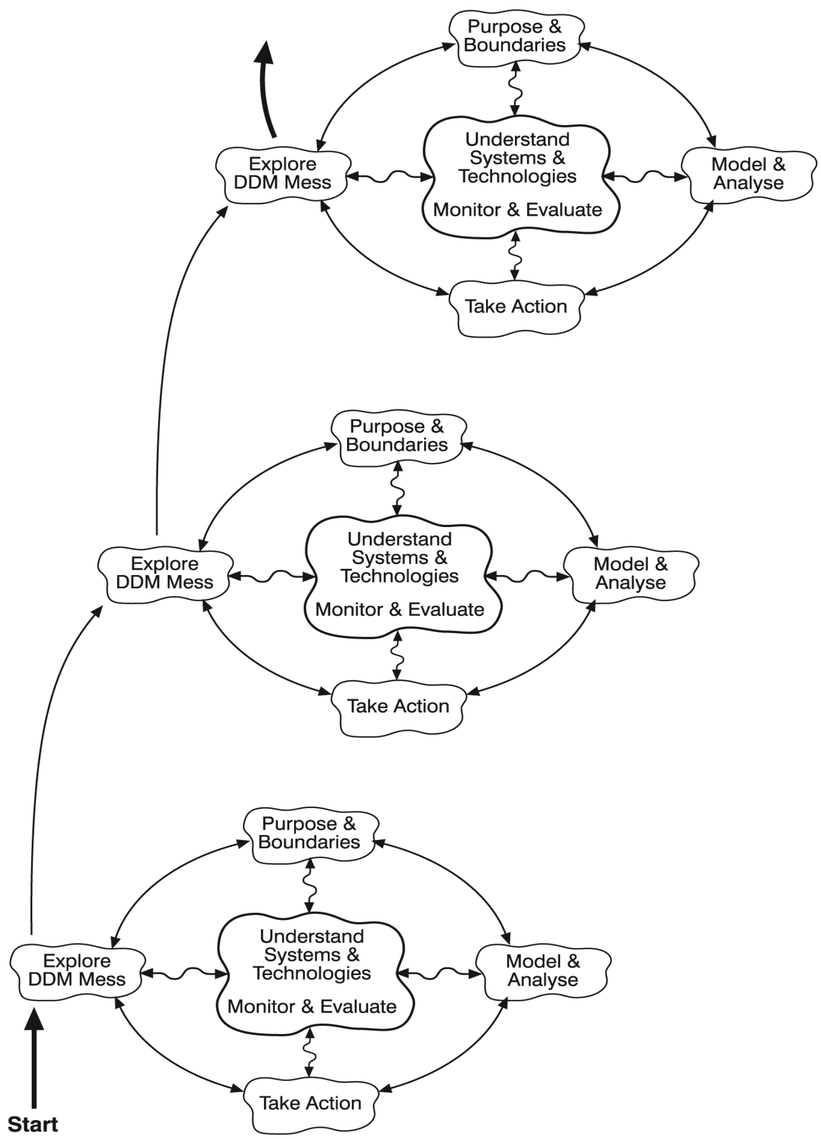

Fig. 2 Cascading iteration through the DDM framework

The idea is to portray the feeling of moving upwards and onwards as the DDM stakeholder progressively works through cycles of the framework.

Whatever the approach, it should be based, in the first instance, on the acquisition of greater under-standing by decision makers, experts and users of systems, technologies and the environment in which these are to be used, to enable the possibility of more successful outcomes.

This paper has not attempted to cover the enormous range of data gathering and analysis, systems diagramming and modelling techniques that are available for making detailed sense of a complex system mess at each stage of the suggested framework. All that is provided here is an overview of how it could be applied to DDM.

Conclusion

The central claims of this paper are that:

1. DDM is messy.

2. The far-reaching implications for commerce and society, of decisions in invisible or opaque specialist fields regarding the regulation and deployment of large information systems, mean they should be matters of concern for all of us.

3. Ordinary citizens working together with experts and regulators will prove more effective than each group acting in isolation. In the application of science and technology to social problems, technologies, systems and policies must be developed together by users and experts. Thinking systemically and considering multiple perspectives will help with this.

4. The default rules of the road in DDM are the laws governing the flow of information and the restrictions built into the architecture of technology.

Some of the decisions we have to make about DDM systems might involve tough new personal, political, regulatory, technical and socio-economic choices, challenges and opportunities. Yet the experience reminds us that, though the context and the tools might change, many of our important guiding principles do not. In the 6th century an Irish monk made a closing argument in the High Court of Ireland relating to a dispute over the copy of a book that could have come straight out of one of the modern DDM cases in intellectual property [Corrigan, 2007, Chapter 1]:

"My friend's claim seeks to apply a worn-out law to a new reality. Books are different to other chattels (possessions) and the law should recognise this. Learned men like us, who have received a new heritage of knowledge through books, have an obligation to spread that knowledge, by copying and distributing those books far and wide. I haven't used up Finnian's book by copying it. He still has the original and that original is none the worse for my having copied it. Nor has it de-creased in value because I made a transcript of it. The knowledge in books should be available to anybody who wants to read them and has the skills or is worthy to do so; and it is wrong to hide such knowledge away or to attempt to extinguish the divine things that books contain. It is wrong to attempt to prevent me or anyone else from copying it or reading it or making multiple copies to disperse throughout the land. In conclusion I submit that it was permissible for me to copy the book because, although I benefited from the hard work involved in the transcription, I gained no worldly profit from the process, I acted for the good of society in general and neither Finnian nor his book was harmed."[18]

In creating, deploying, regulating and operating modern digital communications technologies we have a lot to learn from history.

Notes

[1] Stanford University's Jack Steele Parker Emeritus Professor of International Management.
[2] Both the Germans and the British were adversely myopic about each others' technical capabilities during the War (often to the point of damaging their own war effort).
[3] The UK's Minister of the Interior
[4] *Advocates* by David Pannick (Oxford University Press, 1992) is particularly compelling and lucid on this point.
[5] Professor Emeritus of Sociology, Yale University
[6] The interface between people and machines was a critical factor at Three Mile Island. The control room was large and complex, with a huge number of instruments, controls and indicators but in spite of this would only be staffed by a single operator during normal operations.
[7] One indicator light was covered with a tag label and the indicator relating to the valve was indirect in that it only showed the solenoid (electrically controlled switch) connected to the valve had worked. So although the valve got the signal to close, it stuck open and there was no way for the operators in the control room to know immediately that this had occurred. The label under this indicator light read 'Light On RC RV 2 Open' which would suggest light off meant the valve was closed *not* 'an electrical signal has been sent to the valve telling it to close'.
[8] Though the presidential commission which investigated the incident said there were gaps in the operators' training in dealing with emergencies, all four operators in the control room at the start of the accident had at least five years experience of nuclear reactors in the US Navy before joining the nuclear power generation industry. The US Nuclear Navy created by Admiral Hyman G. Rickover has never had an accident and is renowned for producing top class technical operatives. See *Rickover: The Struggle for Excellence* by Francis Duncan (Naval Institute Press, 2001), *Rickover and the Nuclear Navy: The Discipline of Technology* by Francis Duncan (Naval Institute Press,1990) and *The Rickover Effect: How One Man Made a Difference* by Theodore Rockwell (Naval Institute Press, 1992).
[9] The plant was still a long way from being safe and indeed there was a big release of radioactive gas into the atmosphere two days later. The rest of the story makes fascinating reading, and is told wonderfully well in Chapter 2 of Understanding Systems Failures by Victor Bignell and Joyce Fortune (Manchester University Press, 1984) and many other scholarly tomes. Bignell and Fortune also do a superb analysis of the person–machine interface problems created by the set-up in the control room.
[10] According to Lawrence Mulloy, NASA's Manager of the Solid Rocket Booster project, in his testimony to the Presidential Commission, there were over a thousand "Criticality 1 waivers" on the launch. The failure of any one of these components would have been enough to lead to the loss of the shuttle and the death of the crew.
[11] See also US Office of Strategic Services *Simple Sabotage Field Manual* for US spies, 1944, pp28-32. Declassified in 1976. Available at http://community.e2conf.com/servlet/JiveServlet/download/1090-5-1190/OSS%20Simple%20Sabotage%20Manual.pdf
[12] Sklyarov was detained by the FBI after giving a paper on the nature of the digital locks on Adobe e-books at the DefCon security conference in the US in 2001.
[13] There are some excellent websites on the case at http://www.shirleymckie.com/index.htm and http://www.clpex.com/McKie.htm and the Scottish Parliament's Scottish Criminal Record Office inquiry homepage http://www.scottish.parliament.uk/business/committees/justice1/ScottishCriminalRecordOfficeInquiryHomepage.htm
[14] If we reduce the discount rate to 5%, earnings drop to 95% of what it would be if copyright lasted forever.
[15] See "Performing artists - no longer be the 'poor cousins' of the music business" – Charlie McCreevy at http://europa.eu/rapid/pressReleasesAction.do?reference=IP/08/240&format=HTML&aged=0&language=EN&guiLanguage=en
[16] *T860 Environmental Decision Making: A Systems Approach* (1996–2006) and its successor *T863 Environmental Decision Making: A Systems Approach* (2006-present)
[17] Contact Point

[18] This is a relatively clumsy translation of the original recorded in *Betha Colaim Chille* by Manus O'Donnell (1532). It relies on *Manus O'Donnell, Betha Colaim Chille* (1532) edited and translated by A. O' Kelleher and G. Schoepperle as *Betha Colaim Chille/ Life of Columcille* (University of Illinois Press, 1918), pp178–179, *Saint Columba of Iona* by Lucy Menzies (Llanerch Press, 1992, originally published in 1920), p.25 and *The Legend of Saint Columba* by Padraic Colum (Sheed and Ward, London, 1936), pp76–81.

References

Ackoff, R.L. (1978). *The Art of Problem Solving: Accompanied by Ackoff's Fables*. New York, John Wiley & Sons.

Boyle, J. (1997). *Shamans, Software and Spleens: Law and the Construction of the Information Society.* Cambridge, Harvard University Press.

Boyle, J. (2003). *The Second Enclosure Movement and the Construction of the Public Domain*, 66 Law & Contemp. Probls. 33 (2003), available at: http://www.law.duke.edu/journals/lcp/articles/lcp66dWinterSpring2003p33.htm.

Corrigan, R. (2007). *Digital Decision Making: Back to the Future.* London, Springer-Verlag.

Drahos, P. and Braithwaite, J. (2002) *Information Feudalism: Who Owns the Knowledge Economy*. London, Earthscan.

Franklin, B. (1753). *Method of Deciding Doubtful Matters* Benjamin Franklin (letter to Joseph Priestley)

Higham, C. (2004) *Howard Hughes: The Secret Life*. New York, St. Martin's Griffin & Virgin Books.

Jones, R.V. (1998). *Most Secret War.* London, Wordsworth Editions.

Kretschmer M. et al (June, 2008) *Creativity stifled? Why copyright term extension for sound recordings is a very bad idea*. Bournemouth, Centre for Intellectual Property Policy & Management.

Lessig, L. (1999). *Code and Other Laws of Cyberspace*. New York, Basic Books.

Lessig, L. (2001). *The Future of Ideas: the fate of the commons in a connected world*. New York, Random House.

March, J. (1994). *A Primer on Decision Making: how decisions happen*. New York, The Free Press.

Pannick, D. (1992). *Advocates.* Oxford, Oxford University Press.

Perrow, C. (1999). Normal Accidents: Living with High Risk Technologies. Princeton, Princeton University Press.

Royal Commission on Environmental Pollution 21st report. (1998). *Setting Environmental Standards.* Available from http://www.rcep.org.uk/standards.htm

Saxe, J.G. (1873). The Poems of John Godfrey Saxe, Complete edition. pp 78-79. Boston: James R. Osgood and Company. *The Blind Men and the Elephant* is also available at http://en.wikisource.org/wiki/The_Blindmen_and_the_Elephant

Schneier, B. (2003). *Beyond Fear: Thinking Sensibly About Security in an Uncertain World.* New York, Copernicus Books

Senge, P. (1992). *The Fifth Discipline: Art and Practice of the Learning Organization.* New York, Random House.

Simon, H.A.. (1957). *Models of Man.* New York, John Wiley & Sons.

The President's Commission on the Accident at Three Mile Island. (1979). *The Report of the President's Commission on the Accident at Three Mile Island. The Need for Change: The Legacy of TMI.* U.S. Government Printing Office, Washington, DC.

Vaughan, D. (1997). *The Challenger Launch Decision: Risky Technology, Culture and Deviance at NASA*. Chicago, University of Chicago Press.

Vickers, G. (1970). *Freedom in Rocking a Boat: Changing Values in an Unstable Society.* London, Penguin Books.