When Paradigms Shift: IT Evaluation in a Brave New World

Frank Bannister
Trinity College, Dublin
Frank.Bannister@tcd.ie

Over the years, there have been many foci in the search for IT value. However impending developments in information and other technologies may be about to change the nature of the quest entirely. For example, the prospect of technologically enhanced biological function raises new, difficult and disturbing questions about value that need to be explored. Longer term, developments areas such as cyborg technology, artificial intelligence and robotics could have profound, and potentially disruptive, implications for societies and even humanity as a whole. As of now, there is a rapidly diminishing window of opportunity in which to get our values and value systems clear before a combination of technological advance and market forces overwhelms our ability to make important value choices. This paper explores some of the possibilities that may be coming our way and asks some difficult questions about IT value in what may be a brave new world.

Keywords: IT value, emerging technology, artificial intelligence, robotics, cyborgs, nanotechnology, discontinuity.

1. Introduction: Thinking outside the box

1.1 A short parable

Consider for a moment the electric motor. The electric motor was invented in 1834 by Thomas Davenport, an impoverished, but self-taught, Vermont blacksmith (Wicks, 1999). At the time, the new invention was seen as a possible replacement for other forms of rotational power delivery. Davenport himself saw it as a way of powering the machines in his workshop. Later he came to view it as a possible form of engine for locomotives. Unfortunately for Thomas Davenport, despite patenting his brilliant invention, it did not become commercially successful in his lifetime. In fact its first major commercial success came when somebody had the idea of inverting an electric motor and turning it into a generator. Alas, Davenport never saw this possibility and consequently missed out on making a fortune. He died in 1849, still of modest means.

The relevance of this story in this context is that it illustrates the problem of paradigm limited vision, i.e. people's tendency to overlook the transforming possibilities of new technology. For a long time, the electric motor's potential was considered only in terms of those current technologies that it might replace, i.e. steam, wind and water. Apart from not seeing the potential for electricity generation, nobody at that time dreamed of refrigeration or vacuum cleaners or DVD players. Yet these technologies utterly depend on Davenport's invention. The electric motor was going to change the world – but it was a generation or more before just how much it was going to change things became evident.

This salutary tale informs much of what follows.

1.2 Three categories of evaluation

The time horizons within which IT evaluations are generally discussed, whether this be in the most abstract of theoretical expositions or in the most pragmatic of case studies, fall into three distinct categories:

- First there are studies that focus on the long-term historical economic impact of investments in IT. Brynjolfsson (with Hitt and others) have spent many years exploring the so-called productivity paradox and the cumulative effect of investments in IT on organisations. Brynjolfsson (Brynjolfsson and Hitt 2003) is now sufficiently confident of his findings to pronounce the so-called productivity paradox as near dead as matters and assert that it is now beyond dispute that almost all of the increase in US productivity in the past 20 years is due to IT. Other long term thinkers such as Strassmann (1985) have argued for many years that such effects are only really assessable over even longer periods, maybe as long as half a century.
- Secondly, on a less ambitious scale, there are studies of whether specific investments made over shorter periods have yielded value (or of ways of doing this). Usually such research is in the form of case studies and retrospective analyses. These vary from application of innovative methods to measure value realised to use of well established methodologies such as return on investment, comparison of how different metrics report or combinations of measures (such as the balanced scorecard (adapted from Kaplan and

Reference this paper as:
Bannister, Frank (2005) "When Paradigms shift: IT evaluation in a brave new world" *The Electronic Journal of Information Systems Evaluation*, Vol. 8, Iss. 1, pp 21-30, available online at www.ejise.com

Norton 1992) or the Prudential Appraisal Method (Coleman and Jamieson 1994). Methods can be quite complex as theorists try to distil out the effects of IT from those of other factors and identify the variables that determine the degree of value received. Value for money studies fall into this category.

- Thirdly there is the forward-looking segment of the field. This is concerned with ways of assessing whether or not a potential investment in IT is worthwhile. This is probably the most voluminous part of the literature with, by this stage, dozens of techniques for such evaluations being proposed (and sometimes disposed). In the nature of the task, the horizon here is typically fairly short, usually five to ten years though from time to time studies will contemplate a more distant time horizon. Almost all such studies are at the level of the organisation, be it a firm or a public sector body. Evaluation of impacts at the personal or societal level are relatively rare.

All of these analyses have at least one important thing in common; they are all about *existing* information and communications technologies. These technologies[1] may be well established or they may be cutting edge, but in each case the evaluators or researchers are concerned with either the detectable impact of a known technology or the potential impact of a known, albeit sometimes an emergent technology. Such assessments may be of organisations or, as already noted, less commonly of individuals or economies. In all instances, the technology itself, if not its impact, is understood.

1.3 Time for a re-think?

Ironically, the much debated article *IT Doesn't Matter* (now a book less confrontationally entitled *Does IT Matter?)* by Carr (2003, 2004) reflects this same mode of thought. Carr argues that, from a business perspective, there is not much exciting left to come from information technology. It may be possible to build more efficient supply chains, make processes even faster, improve customer relationship management or even find out how to manage knowledge better using machines rather than people or organisations, but if everybody is doing these things, where is the strategic advantage? Carr's thesis *per se* is not the issue in this paper. What is of interest is the nature of some of the attacks on it from scholars and professionals. Several of these argue that Carr's vision is hopelessly limited; that he does not understand the nature of technology nor does he take sufficient account of what is to come (see Stewart *et al*, 2003). An uncomfortable question that might also be asked is can the same criticism be made of IT evaluation? Is it time for IT evaluation researchers to start thinking about what is to come, i.e. outside the box?

It is, of course, well established that technology can have unexpected effects on both individuals and organisations. The way people use a technology may not be that for which it was designed, even where the mechanics of the technology itself are known and well understood. A simple example is the explosion in the use of mobile telephones for text messaging – something not anticipated by the telecoms companies. Another example comes from the early 1980s when there was great excitement about the 'home' PC and all the wonderful things for which people would use them. In practice, those that did not end up gathering dust, ended up being used for playing computer games. The ability of pundits, be they academics or business gurus, to forecast the impact of a technology is far from infallible. If this is true of established technology, how much more true it is likely to be of technology that is not yet with us? It is the conjecture of this paper that there are impending developments in technology for which the weapons currently stockpiled in the arsenal of evaluation techniques are not only inadequate, but are quite inappropriate. This is a bold conjecture and not easy to establish. Nonetheless, this paper will at least try to make the case that there may be a case. In so doing, it will ask some fundamental questions about the nature of evaluation itself in the context of paradigm shifts.

This will be done in two steps. First, although to support the argument it is not necessary to speculate on exactly what the future will be, it is necessary to look at some of the ways that potential developments in information technology could change individuals, organisations and societies over the next few decades. Five probable (or at least possible) developments will be considered. Secondly, it will be argued that contemporary evaluation tools either do not work at all with such developments, or at least do not work very well. Finally there are some reflections on what this might mean in terms of new opportunities for IT evaluation.

[1] The word 'technologies' is used here and elsewhere to mean information technologies unless explicitly stated to the contrary.

2. Brave new worlds

2.1 Knee Points

"*Prediction*", said Yogi Berra, "*is very hard, particularly when it's about the future*". When surveying the world of IT futures, it is not practical to pursue every possible avenue, so what follows is necessarily selective. There are many technologies which have the ability to radically alter the way people live and work, organisations operate and societies function. Many of these are closely interrelated; indeed all are interrelated at some point and it is frequently at the intersection of technologies that the most challenging issues arise. This section describes some of the more immanent possibilities.

An important point to bear in mind is that some of these technologies are approaching what Kurtzweil (1999) calls the 'knee of the curve'. By this he means that point on an exponential growth curve at which a technology which has been growing slowly for some time suddenly takes off. The concept is illustrated in figure 1.

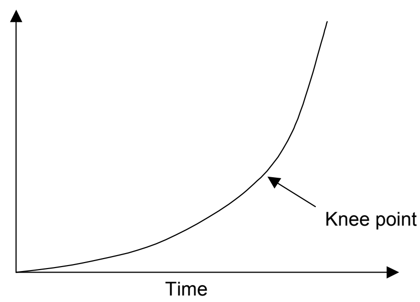

Figure 1: Knee point of a curve

One of the characteristics of a knee point is that once a technology passes it, it becomes exceedingly difficult to control it thereafter. The technology takes on a life of its own and subsequent social or even legal constraint may be impossible. Beyond this point the technology tends, as it were, to find its own level[2]. Whether social control is necessary or even desirable is a debate for another day (for those interested, McKibben (2004) discusses this issue at some length). The point to note here is that the IT/IS evaluation community may be faced with tricky questions about several of these technologies sooner rather than later.

2.2 Artificial Intelligence

Artificial intelligence is a broad church, encompassing technologies from rule based programming to image recognition. Some of these technologies are now at the knee of the curve. Some are still some distance away and yet others would still appear to be in the realms of science fiction though the latter have an increasingly disconcerting habit of become fact. While the areas in which AI is likely to have a radical impact over the next decade are robotics, control systems and localised intelligence/interfaces, the more remote possibilities of consciousness and self awareness need to be considered. A key step up on the AI tree is the adapting/learning system. At the most basic of levels this is now commonplace. Contemporary developments in robotics are one aspect of this. Anybody who has used a voice recognition system will be familiar with the concept.

For the purposes of this discussion, robots will be defined as differing from (mere) automation in two ways. First, within limits, they are capable of being programmed to do a wide range of *non predetermined* tasks. Secondly, robot systems have a capability to learn and more advanced systems can adapt to the problems they face and/or the environment that they are in. As a technology, robotics (as opposed to automation) has been slower to take off than many of its more misty eyed advocates anticipated although the market for robotics is forecast to rise to $16 billion by 2007 (Quality 2003). Despite major investments in robotics, today's factories are still extensively dependent on human labour. Where humans have been displaced in the workplace or in the home it has tended to be by automation rather than by adaptive machines. Nonetheless the prognosis must be that robots will eventually be able to do many of the repetitive manual tasks currently done by humans. This may range from fitting a door on a car to serving a burger and fries.

Localised intelligence/interfaces are a more visible (or audible) development in AI. Such systems are already partially embedded in applications such as directory enquiries systems and help lines. When a customer dials a directory enquiry number today, there is a reasonable probability that she will be answered by a computer. Already experiments have taken place with fully automated telephone booking systems and help desks. Current research in Interactive Voice Recognition (IVR) includes emotion recognition so that machines can respond to users more

[2] It is easy to think of examples of this phenomenon. A pertinent one is the use of performance enhancing drugs in the Olympics. Another example is peer-to-peer sharing of music files.

appropriately (Yacoub et al 2003). Such systems try to solve the problem themselves and only pass it to a human operator if they cannot resolve it or for the key decision(s). There is no reason why such systems should not be widespread by the end of the decade even at the level of domestic appliances or office equipment.

The ultimate goal of AI is the self aware or conscious machine. Currently this is still science fiction, the world of the film *The Terminator* or HAL in *2001*. Whether a self aware, thinking computer is possible is hotly debated. Popular authors such as Kurzweil (1999) have long forecast this development. The late Nobel laureate, Francis Crick spent much of his later life trying to establish that the brain was nothing more than a biological machine (Crick and Koch 2003) and many of his fellow life-scientists hold similar views. One of the most noted proponents of this theory in the IT world is Marvin Minsky (1988). Philosophers and others have challenged the concept, one of the most forceful being Penrose (1989, 1995). However, if such a development were to be possible, the evaluation of the consequences would raise profound issues including, *inter alia*, for evaluation.

2.3 Cyborg technology

If AI is about machines behaving like humans, cyborg technology is about machines merging with humans[3]. This is the world of film *The Matrix* and the television series *The Six Million Dollar Man*. One of the most passionate advocates of this type of technology is Warwick (2003) who has gone so far as to have a chip implanted in his arm[4]. This chip enabled devices in Professor Warwick's laboratory to detect his presence. Doors would open on his approach and greet him. Chip and bolus technology has been the subject of experiments with animals for tracing purposes (where it causes problems, not least from the tendency of such chips to migrate within the animal). In the Netherlands, customers at the Baja Beech club in Rotterdam have chips embedded in their upper arm in order to get automatic access to the VIP area of the club[5]. Humans already walk around with artificial limbs and pacemakers; having intelligent devices on board is only the next logical step.

Putting Radio Frequency Identification (RFID) chips into one's upper arm is one the simplest applications for this type of technology. The ultimate goal is to link the body's own central processing unit, the brain, to the computer. In theory, there are many ways in which this can be done. A crude method is to provide a link to external machines via electrodes attached to the skull. As fiction, this was the theme of Craig Thomas' book *Firefox* (1990) in which the Russians had supposedly developed a fighter aircraft where the pilot's brain was directly connected (via external electrodes in his helmet) into the aircraft's control system. To fire a weapon, for example, the pilot merely needed to think of the required action. Enormous research is going into making this particular piece of fiction into fact. With fighter aircraft the race may be between this technology and aircraft that have no pilot at all. Warwick (2004b) describes just such an experiment. A further stage would be to implant such a communications system *within* the brain so that people could connect to a control system or the Internet via (say) WiFi anytime they desired. Imagine a house that responded automatically to one's wishes, that would sense one was cold and turn up the central heating or automatically find the TV programme you wanted to watch or music you wanted to hear. In a business context, such a system would enable employees in an organisation to be in constant contact with each other without the hassle of bulky laptops or vulnerable mobile phones. In the longer term there is the possibility of clustering human brains in the way computers are clustered today. It would certainly give a whole new meaning to the concept of teleconferencing.

The ultimate development of cyborg technology would be with the implantation of enhancing technology within the body itself. To take a silly, but not entirely implausible' example many people are poor at mental arithmetic. Suppose a company developed an arithmetic processor that could be connected to the brain in such a way that the user could pass any arithmetic problem to it and get an answer[6]. Such an enhancement could even be biological in nature (see below). Or what about the ability, as in the film *The Matrix*, to program in any mental skill[7]? Or how about a memory enhancing device? Many of the science fiction features first seen in the television series, *The Six Million Dollar Man* in 1975 are today close

[3] Also a common theme in science fiction (e.g. the Borg in the *Star Trek* series or the film *Demon Seed*).
[4] www.wordiq.com/definition/Kevin_Warwick
[5] The author thanks Professor Egon Berghout for drawing his attention to this interesting example.

[6] Readers who think this is far-fetched might wish to look at Warwick (2004a).
[7] Programming in physical skills presents more fundamental problems!

to becoming realities. However even the bionic man was confined to increased physical powers; increased mental powers are a different matter entirely.

2.4 Miniaturisation

Another potential source of change, or it may be more accurate to say a catalyst for change, comes from developments in areas such as nanotechnology, biotechnology and quantum computing. At the moment, the first of these is nearer to large scale practical application (for IT) than either of the latter. Nonetheless the possibilities for radical change here are formidable. Biologically based computing, for example, offers the potential to put vast amounts of processing power into the tiniest spaces – a few spoonfuls of water according to one group of Israeli researchers[8]. Quantum computing offers even more staggering possibilities in terms of power per cubic nanometre. As a result, the potential exists to build in, say, a high level of intelligence into any machine or device not to mention into humans. There are also implications in, for example, cryptography where quantum computing could lead to unbreakable codes (Singh 1999).

Miniaturisation is not, in itself, that significant from an evaluation perspective. After all, the world has lived with Moore's law for three decades now. It is when miniaturisation is combined with developments in AI and cyborg technology that is acquires critical significance. To take but one aspect of this, contemporary surveillance systems still have a limited ability to process the amount of information they gather. The ECHELON system may be able to monitor most of the world e-mail, but it relies on slow and unreliable keyword searching to deduce information (Bamford 2001, Levy 2001). A quantum computing system could read every e-mail, listen to every mobile telephone conversation and digest the contents from every CCTV camera on the planet – a thought that gives a whole new meaning to the concept of a surveillance society. A particular case of miniaturisation is the concept of nanobots. Applications for nanobots are legion from cleaning up waste to keyhole surgery. Most of these applications are benign, but there are risks of undesirable consequences, in particular in control of such technology and in undesirable uses.

[8] http://search.eurekalert.org/e3/query.html?col=ev3rel&qc=ev3rel&qt=drop+of+water

2.5 Networking

In many ways, networking is the most prosaic of trends. Unlike some of the technologies discussed in the preceding sections, this is a here and now. The means to connect and collect is already developed if not yet widely deployed. Although, unlike some of the other technologies discussed, the impact of networks is moderated by basic logistical constraints, when combined with artificial intelligence and miniaturisation then the potential effects are profound.

For example, in theory networks could connect everybody in the world or in a state to a central 'nervous system' of some sort. A crude version of this already occurs with the telephone system. Mobile phone systems mean that people can be contacted any time and (provided they leave their phone on) their movements traced. Already some nervous wealthy parents are putting RFID chips into their children so that they can trace them if they get lost or are kidnapped. Tagging of criminals is common. It is not difficult to envisage a world where the whereabouts of anybody is knowable at any time, and where everybody who does not make energetic efforts to avoid it is linked into some soft of universal communications matrix. There are considerable attractions in this for some people. The fact that people are prepared to leave their mobile phones on in place as diverse as bathrooms and lecture theatres suggests that many of us are all too willing to be on-call 24 hours a day. From an evaluation perspective, this presents quite complex challenges. What of the risks? What if a virus got into such a network, a possibility explored in fiction by Stephenson (1993). Like the other technologies, networks present opportunities and risks and it is this that causes problems for evaluators.

2.6 Electronic identification

A related issue electronic identification. Reference has already been made to the surveillance society. However there is a broader issue here than a simple fear of big brother looking over everybody's shoulder. It is will shortly be within the reach of technology to issue a citizen identity card with unique identifying information such as a person's DNA or retina pattern to all citizens. Identity cards have been a fact of life in many countries democratic as well as undemocratic for decades. Some countries, including the UK and Ireland do not have a citizen identity card although the UK government is keen to introduce one (Stephens 2004).

There are obvious benefits from a universal identification system from countering terrorism to preventing fraud. Apart from the implications for privacy and civil liberties, there are also risks from identity theft. These risks differ by an order of magnitude when an identity system becomes electronic. In a society that is ever more dependent on carrying an increasing volume of personal identification, the downside risks to both citizens and society increase. To make matters more complicated, questions of freedom and the nature belonging to a society means are made more complicated. Can somebody remain a citizen without having a card? What is the position of those who refuse to comply? For evaluators, these questions raise difficult issues with soft benefits and disbenefits which have, to date, seen little discussion within the community.

3. Some problems in evaluation

3.1 Conventional methods

Traditional or conventional evaluation of IT assumes that IT has a cost (which may or may not be exclusively financial) and benefit(s). The central question in most of the literature is how to measure these benefits (although as Bannister *et al* (2003) have argued, measuring the cost is not that straightforward either). Benefits can be individual, organisational, economic or social or any combination of these, but, and it is an important 'but', the nature of the change wrought neither threatens the whole basis of society nor our social understanding nor our understanding of ourselves. Each of these is relevant in the context of potential future developments. If the evaluation community has struggled with valuing the impact of current systems and technologies, then the problems in evaluation discontinuous technologies are likely to present an even more formidable challenge.

3.2 Buying some brainpower

To illustrate the problems it is easiest to use examples. Taking cyborg technology as an example, consider the following thought experiment.

A computer company produces an electronic device that can be connected directly to a human brain in order to enhance a person's mathematical ability. John, a dealer in the financial services industry, is contemplating having such a device installed. It is quite expensive, costing, say, €100,000 including implantation and after sales support. There are no known risks from this technology and running costs are minimal although its estimated working life is 20 years. How might this be evaluated?

In this instance, the financial cost is clear. What about the benefits? A crude approach might to be assess the increase in John's lifetime earning power or by (say) his ability to make a killing on the futures exchange through a capacity to compute arbitrage rates faster than the market. Adopting this worldview, standard return on investment techniques can be used and value can be quickly assessed. There might also be slightly more subtle benefits from enhancements to John's lifestyle. He may be able to impress his friends with his mathematical skill or keep track of what he spends in a supermarket as he fills up his trolley. The feeling of well being from this might be worth something although adaptation (Schwartz 2004) makes it likely that this effect will sooner or later wear off. Furthermore, the pleasure might be diluted by the vague feeling that he is cheating, that this ability is not really 'his', but that of a device to which he has access.

Another question is who pays the €100,000? If John pays it from his own resources, then there is one kind of evaluation problem. If his employer pays it, then the issues are different[9].

3.3 On being human

All of the above questions are easily to deal with when compared to the problems raised by the following awkward question: *is John the same person that he was before the implant?*

It could be argued 'yes he is' if one follows the line that John's improved ability is not really 'his' at all. Apart from integration and speed, there is no essential difference between John and somebody with a good laptop and fast fingers. John's integrity as an individual is not therefore compromised. On the other hand, it can also be argued that there *is* a difference between such an implant and a laptop. The laptop is out there. Anybody can use it. It can be switched off. The implant is uniquely attached to John, it may even be personalised or tailored to his physiology, and the question '*does John still exist apart from the implant?*'

[9] In answer to the obvious question why would an employer pay for such a device, the short answer is that employers currently pay to improve their employees in various ways from putting them on training courses to giving them free medical check-ups.

especially if the implant were biological, is a valid one. Furthermore, would John be able to differentiate between the implant and his 'normal' brain. And if one answers no to that question, then some part of John has changed: certainly his ability has been enhanced, but his integrity as an individual has been diminished or at least altered by the fact that he is now partially constructed. The question then is, what is the value of this change/loss?

To make this case more vivid, suppose that instead of a computational implant, John has a memory implant. As the computational implant contains powers John does not have, so the memory implant can contain information that John has not 'learned'. It might contain a dictionary or an encyclopaedia or two or a language module. And how about some pre-programmed happy memories? It is alleged (it is a much debated topic) that the brain can generate its own false memory (Stanton 1997); imagine what it could do with a little technical help. This raises uncomfortable questions about the meaning of reality, at least from John's perspective. He may be unable to distinguish his real past from an artificial past created for him. Again this issue has been explored in science fiction. One of the most famous examples is by Dick (2002)[10]. It needs to be explored in the evaluation community.

The question that arises in both these scenarios of what it means to be human? Is John more or less human because of these enhancements, or is the nature of his humanity unaffected. He now has increased powers, but less integrity, because he is now, in part, designed by somebody else. However John also has an artificial hip and two dental crowns. Is the brain implant qualitatively different? Furthermore, his perception of reality is in part constructed by outsiders and is not part of his authentic experience. Pushing this to extremes, it is possible to envisage a whole range of IT-powered enhancements that might improve John's 'performance' at the progressive expensive of altering who John is. Is this a cost and if so should it be added to the financial cost of the implant? Even the counter argument that as long as his personality is unchanged, John is unchanged, does not stand up to close scrutiny. How is it possible to disconnect personality or behaviour from memory for example? Finally, how are such benefits and 'costs' to be evaluated?

[10] This short story was the basis of the film 'Total Recall'.

3.4 What is a life?

Another type of problem (if not several such problems) is/are posed by AI. To take an extreme case, suppose that a self-aware, conscious machine is developed. By any definition this would be a new life form. Furthermore, given the dynamics of the situation, unless its design specifically inhibited it, it is likely to be a form of intelligence that will rapidly surpass the mental capacity of its inventors (provided that is they keep supplying it with the extra processing power and memory it needs) [11]. The impact of such a development on humanity, not to mention business, is hard to assess. Like cloning and genetic engineering, a conscious machine may subtly (or maybe not so subtly) change the meaning of what it is to be human or at least our understanding of ourselves. This poses major problems for evaluation of such technology. How would one value such a creation? Philosophers, politicians and insurance companies debate the value of a human life. What value would one put on such a machine life (even if it did not have any feelings)? Even avoiding this question, what are the business implications? Such a machine might soon make most managers redundant. This may be cost effective, but the implications for those making the investment could be redundancy and given that many large corporations these days seem to be run mostly for the benefit of managers (Galbraith 2004), who would want to make such a decision?

3.5 Joining the matrix

Another example is the potential of new technologies to disrupt society in too short a time for humanity, be it individuals, organisations or societies to react, is the ability of computers to intrude into people lives. One of the most difficult of evaluations is the trade off between personal security and privacy. With modern technology, it is possible to make people's lives more secure in a variety of ways. Identity cards with biometric information, closed circuit television, RFID tagging, mobile phone tracking and other technologies can be employed to ensure that citizens are not defrauded, mugged, kidnapped or lost. But this is achieve at a price to privacy and a risk of misuse by private organisation or by the state.

Balancing risks in this situation is not simple and again represents a challenge to the evaluation community. A large part of the

[11] Dick (1968) also wrote a famous book on this subject. It was made into the film 'Blade Runner'. In the story, the lifespan of androids is deliberately kept short.

problem in evaluation is the asymmetry in risk perception. People perceive the risks from, say, injury in a terrorist attack to be much higher than it is in practice. In parallel, they fail to perceive potential risks to freedom, privacy and even democracy from technologies ostensibly designed to prevent terrorist attacks. Recently two US commentators, Dash (2004) and Rosen (2004) have looked at this problem from a broad perspective. However it is also an IT evaluation problem and one which, to date, the evaluation community has given little or no attention.

3.6 Problems for evaluators

The problem for the evaluator of the type of technological developments outlined above is which evaluation techniques are appropriate? Financial and economic methods hardly seem appropriate. Consumer satisfaction, organisational improvement, information economics, balanced scorecards; none of these seem adequate to cope with the philosophical problems raised. Questions of risk, judgement, bias, uncertainty, humanity, meaning and society are all impacted by these developments. The English poet John Donne once wrote that '*no man is an island entire of itselfe'* and this has never been more true than it is today. An investment by any business in these technologies affects far more than the business. Evaluation cannot be ring fenced in these circumstances.

The more transformative an IT induced change, the more difficult is can be to evaluate. There are several reasons for this. First, as was pointed out above, the metrics applicable in the before and after situation may not be the same because, to use a well known phrase, apples are not being compared with apples. This may not be a problem when looking at the investment in retrospect, but it is a definite difficulty in *ex ante* evaluation. The problem here is that the type of change that might happen with these technologies is more than merely transformative, it is discontinuous. It represents a potential radical departure from the known.

Secondly, the more transformative the change, the more likely it is that evaluation will depend on subjective metrics, i.e. metrics which depend on the views or judgments of actors and in particular 'experts' and so-called gurus. It is clear that such views and judgments are not just problematic from a psychological and judgemental perspective, but are time dependent, so the impact may depend on when the measurement is taken (Schwartz 2004, Myers 2002).

Thirdly, the more transformative a technology, the greater the degree of uncertainty about the outcome and the greater the scope for errors of judgment. The law of unintended consequences will inevitably apply. There are many cases of IT investments where the result has not been quite what was expected[12]. In evaluating technologies such as AI and cyborg technology, IT evaluation may therefore have to make much more extensive use than heretofore of general risk evaluation tools such as scenario analysis (see, for example, Wright 2001) and sophisticated risk analyses methods.

In summary, the challenge presented by some of the technologies now evolving is that they are more than radically transformative, they are discontinuous technologies. The evaluation toolset currently available is not capable of providing meaningful assessments of such technologies. It is necessary to reach not just for new methods, but for deeper philosophical tools.

4. Some reflections

In his keynote speech at the 10th European Conference on IT Evaluation (ECITE) in Madrid in 2003, John Ward, reviewing the state of IS evaluation ten years after the first ECITE, suggested that maybe after three decades of academic attention, the topic of IT evaluation was running out of steam. There were many hundreds of papers, books and articles now in print on the topic, but there was still no one agreed approach or agenda. Instead there was a large toolkit of techniques, none of which was entirely satisfactory and which users therefore mixed and matched as the circumstances required.

A more accurate reflection (as a glance of the proceedings of that particular conference shows) would have been that there is no shortage of new ideas, interesting case studies, new applications and new ways of combining techniques around. Research and thinking in the field remains healthy and active; special issues of journals on the topic of IS evaluation are still being published. But it is probably accurate to say that there have been no big conceptual breakthroughs in quite some time (although there are some approaches that may hold promise, for example Halpin and

[12] For an amusing, but insightful, collection of examples see O'Boyle (2000) or Harvey-Jones and Tibballs (1999)

Stapleton's (2003) application of complexity theory to post implementation evaluation). Pluralism is increasingly the name of the game. Approaches to evaluation tend to be holistic and reflective (Cronk's (1999) concept of 'holistic construal' is a good example of this school of though) of the complex realities that are involved in all but the most simple of situations. A good example of this type of rounded approach can be found in Curley (2004).

This paper proposes a radical new line of enquiry for evaluation of IT, one that looks beyond the known and the short term future and contemplates the impacts of discontinuous technology. In so doing, this paper has raised questions rather than attempt to provide answers. One possible way forward is for the world of IT evaluation to engage much more with current thinking in decision sciences and in risk analysis. Nonetheless there remains a major intellectual challenge in evaluating discontinuous developments. Evaluation here has to move beyond the financial and economic, beyond the conceptual toolkit of the current literature and even beyond questions of risk and uncertainly into fundamental questions about the nature of organisations, humanity, meaning and society. For a subject that is engaged with technology, IT evaluation has tended to be fairly tame in its remit[13]. Dealing with the types of challenge and potential for discontinuous change which IT may present over the next few decades will require new tools and engagement with philosophers from Wittgenstein to Midgley who have wrestled with this problem. What may emerge from this may be quite different from the type of discourse that has dominated the field since these issues were first broached in the 1960s. To paraphrase another piece of science fiction, it may be evaluation Jim, but not as we know it.

References

Bamford, J. (2001) Body of Secrets: How America's NSA and the UK's GCHQ eavesdrop on the world, Century Books, London.

Bannister, F., McCabe P. and Remenyi, D. (2003) "IS costing: the case for a reference model", Southern African Business Review, 7, 1, 1-16.

Brynjolfsson, E. and Hitt, L. (2003) "Computing Productivity: Firm-Level Evidence", MIT Sloan Working Paper No. 4210-01, June 2003.

Carr, N. (2003) "IT Doesn't Matter", Harvard Business Review, May 2003, 41-49

Carr, N. (2004) Does IT Matter?, HBS Press, USA.

Coleman, T. and M. Jamieson (1994) "Beyond return on investment: evaluating ALL the benefits of information technology" in Willcocks, L. (Ed.), Information Management, The evaluation of information systems investments, Chapman Hall, London, pp 189-206.

Crick, F and Koch, C. (2003) "A Framework for Consciousness", Nature Neuroscience, 6. 2, February 119-126.

Cronk, M. (1999) "Understanding Complex Information System Constructs through Holistic Construal" in Brown A. and D. Remenyi (Eds.), Proceedings of ECITE 1999, Trinity College, Dublin, pp 69-76.

Curley, M. (2004) Managing Information Technology for Business Value: Practical Strategies for It and Business Managers, Intel Press.

Dash, S. (2004) The Intruders: Unreasonable Searches and Seizures from King John to John Ashcroft, Rutgers University Press, New Jersey.

Dick, P.K. (1968) Do Androids Dream of Electric Sheep?, Del Ray Books

Dick, P.K. (2002) "We can remember it for you wholesale", in Selected Stories of Philip K. Dick - Volume 2, Pantheon Books.

Galbraith, J.K. (2004) The Economics of Innocent Fraud, Houghton Miflin, Boston, MA.

Halpin, L. and Stapleton, L. (2003) "A theoretical framework based on complexity theory for evaluation of large-scale Information Systems Development projects", in Remenyi, D. (Ed.), Proceedings of the 10th European Conference on IT Evaluation, MCIL, Reading, UK, 297-306.

Harvey-Jones, J. and Tibballs, G. (1999) Business Blunders, Constable and Robinson, London.

Kaplan, R.S and D.P. Norton (1992) "The balanced scorecard - measures that drive performance", Harvard Business Review, Jan-Feb, 70, pp 71-19.

Kurzweil, R. (1999) The Age of Spiritual Machines: When Computers Exceed Human Intelligence, Penguin Books, New York.

[13] There are, of course, some exceptions to this general rule!

Levy, S. (2001) Crypto, Viking (Penguin books), Harmondsworth, UK.
McKibben, B. (2004) Enough, Staying human in an engineered age, Owl Books, New York.
Minsky, M. (1988) Society of Mind, Simon and Schuster, New York.
Myers, D. (2002), Intuition: Its powers and perils, Yale University Press, New Haven, CT.
O'Boyle, J. (2000) Wrong: The biggest mistakes and miscalculations made by people who should have known better, O'Mara Books, London
Penrose, R. (1989) The Emperor's New Mind, Oxford University Press, Oxford.
Penrose, R. (1995) Shadows of the Mind, Oxford University Press, Oxford.
Quality (2003) Robotics Market to Double by 2007, 42, 9, p15.
Rosen, J. (2004) The Naked Crowd: Reclaiming Security and Freedom in an Anxious Age, Random House, New York.
Schwartz, B. (2004) The Paradox of Choice, Harpur Collins, New York.
Singh, S. (1999) The Code Book: The science of secrecy from ancient Egypt to quantum cryptography, Fourth Estate, London.
Stanton, M. (1997) "U-Turn on Memory Lane", Columbia Journalism Review, July/August, 36,2, 44-49.
Strassmann, P. (1985) Information Payoff: the Transformation of work in an electronic age, Free Press, New York.
Stephens, P. (2004), " A prime minister with no regard for liberty", Financial Times, December 21st 2004.
Stephenson, N. (1993) Snowcrash, Penguin Books, New York.
Stewart, T. and others (2003) Do IT Matter? An HBR Debate. Letters to the Editor, Harvard Business Review, June, 1-17.
Thomas, C. (1990) Firefox, Harper Paperbacks, UK.
Warwick, K. (2004a) "Linking Human and Machine Brains - Why you should volunteer", Proceedings of the 5th International Conference on Creative Thinking, ed. S.Dingli, Malta University Press, to appear (2004)
Warwick, K. (2004b) "Mind Blending", People Management, 10, 7, 32-33.
Warwick, K. (2003) "Cyborg morals, cyborg values, cyborg ethics", Ethics and Information Technology, 5, 3, 131-137.
Wicks, F. (1999) "The Blacksmith's Motor", Mechanical Engineering, July.
Wright, G. (2001) Strategic Decision Making: A Best Practice Blueprint, John Wiley & Son, UK.
Yacoub, S., S. Simske, X. Lin and J. Burns (2003) "Recognition of Emotions in Interactive Voice Response Systems" paper presented at Eurospeech 2003, 8th European Conference on Speech Communication and Technology, 1-4 September 2003, Geneva, Switzerland. Available at www.hpl.hp.com/techreports/2003/HPL-2003-136.pdf.