19

Ethical imperatives in Decision Support Systems design

Andrew McCosh
Department of Business Studies, The University of Edinburgh, William Robertson Building, 50 George Square, Edinburgh EH8 9JY Scotland. Tel: 0131 650 3801, fax: 0131 668 3053. E mail: a.mccosh@ed.ac.uk

Abstract

The paper addresses the need to build ethics into the process of designing decision support systems. Two examples are offered, in which the essential concept of a DSS is changed considerably when the priniciples of ethics are applied, taking the decision models far beyond the basic quantitative and economic variables customarily employed. Ethical concepts are described first, and their applicability to DSS design is then demonstrated.

Keywords

Ethics, Decision Support Systems, Investment Appraisal Theory

1   INTRODUCTION

The tradition in DSS design work has always been to ensure that the quantitative aspects of the problems are handled with care and correctness, and to ensure simultaneously that the user groups and other affected personnel within the organisation are satisfied by the services which it enables them to provide.   This paper contends that there is need to go beyond this. In some instances at least, it is necessary to go beyond getting the sums right, and beyond ensuring the supportiveness of the insiders who have to work it, to the stage of establishing that the system gives answers which are morally sound.

2   WHAT IS A SYSTEM?

The word system has had the misfortune to become fashionable, and it is therefore used to mean all sorts of things. The word is in decline as a precise scientific concept. This is a shame; it is a very useful word to describe certain specific combinations of people and of artefacts. I shall define it below; this definition is very similar to other cybernetic definitions.

A system has a series of attributes. These include (1) its membership list, (2) its purposes or functionality, (3) its self-controlling procedures, and (4) its mechanisms for finding out what is going on around it. We shall look at each of these in turn.

A system has members, and therefore also has non-members; it is possible, at any given moment, to list the members. There may be a mechanism for new members to join, and there may be a mechanism for existing members to depart, but this is not essential in all systems. In the present paper I shall be confining myself to that subset of systems in which some of the members are human: at the same time, it should be appreciated that the theory of systems does not require this attribute.

A system has functionality. A system is created in order to fulfil some need. Hobbes has shown (Hobbes, 1898, and also Kavka, 1986) how primitive humans in a 'state of nature' create a communal system for their own protection and to reduce the number of people they have to fight. The Apollo system, which had computer, rockets, and people as principal components, was created to land a man on the moon. A church (in the sense of a denomination) is created to enable like-thinking worshippers to promote their views and to worship the way they want to. It is sometimes convenient to look at a church as a system. A market is certainly a system, and each market is itself a component in a larger system of interacting markets. The functionality of a market is to enable people to enhance their material welfare by gaining access to goods and services they cannot make (or make as well) themselves. In each case, there is a goal to be fulfilled.

It is not unusual for a system's functionality to change, especially when it is new. There may be an initial formal goal for the system, which moved people to set it up, and perhaps to pay for setting it up. There may be a significantly different official or formal ultimate goal for the system, which is what it is officially trying to do when it reaches its stable state. Almost certainly, there will be an array of by-product consequences of the system's operations, which may be helpful to the fulfilment of its main goals, may hinder them, or may totally obstruct.

The World Bank and IMF, for instance, were set up by men of the highest integrity and the best of intentions, to assist less developed nations in catching up economically with the remainder. It actually became an economic 'body shop', into which wrecked economies were towed, to be hammered back into something resembling the right shape. This process works quite well for advanced and basically sound 'Mack Truck' economies, which have got into temporary difficulties. The hammering process is not so appropriate for inherently weak 'Trabant' economies, which just cannot absorb the structural adjustment blows. The hammerees generally are unhappy about this process, mainly because they had previously had a perfectly workable (albeit horsedrawn) economy before the salesman came along with the new supercharged model, but failed to tell them it could not be expected to stay on the road with a learner driver. The stated goal was to push progress into locations where market forces

would not have pushed it, and to promote growth where the markets would not have promoted it. It appears the market was right, in that the side effects have (in some instances) caused more damage than the intended goals have caused benefit.

It is important to note that the functionality of a system is not a permanent and unchanging one. A true cybernetic system contains, within itself, procedures for monitoring how it is getting along, and how far it is fulfilling its function. The monitoring subsystem will have the capacity to make adjustments to the way the system operates. Above the monitoring system, there will be a metasystem which monitors the viability of the present set of objectives, with powers to amend these if need be.

This hierarchy of control systems can be carried to any level one may desire, but it is unusual for any system containing humans to have less than three to five layers, and it is similarly unusual for a system without humans as members to have more than two or three layers. (Beer, 1979)

In every important instance, there is a cyclical process in a system. The activity of the system takes place. The amount that has been achieved is appraised. The monitoring system considers whether this is good enough, or whether changes are needed. The changes, if any, are put into position. The activity of the system takes place again. And so it goes on.

The sensors through which a system checks its environment are the fourth set of attributes mentioned above. These may perhaps be best illustrated by a conspicuous failure. The IBM company was so hierarchical, so ponderous, so inwardly focused, and so committee-bound that it took twelve years after it had successfully entered the microcomputer market before it worked out that these little but powerful devices would affect its mainframe business. They are only just beginning to recover from this severe sensor failure now.

A number of important theories have been developed about systems. The most important, at least for the present purpose, are those developed by Ashby (1957) and by Beer. Ashby, especially, has shown that the system you need to control another system has to be at least as complex as the system it is supposed to control. This point is vitally important, as we shall see below.

When we discuss a Decision Support System, I take that to be a special case of the general man-machine system discussed above. The combination of the people and the machinery has functionality, as all systems have, but this functionality is especially designed to enable a particular class of decisions to be taken. The literature on DSS design and development is well developed, and there are many effective models for the design process and for the proper implementation of the resulting systems. The papers by Ferioli and Migliarese, by Berztiss, and by Brezillon and Pomerol, all in this book, are good examples of the processes involved. The mission of the present paper is to suggest that, for some situations, it is not quite enough to tackle the problem this way. We need, in some but not all cases, to take formal account of ethical issues if we are to provide a genuine and enduring support to the decision makers we are seeking to serve and to the publics which they, in turn, are trying to help.

In most instances, it will be shown that the ethical dimension of DSS design and implementation will involve us in a deeper search for the people who will be affected by the decision we are about to construct and to implement. It is not enough just to take account of the manifestly affected. We must consider the people, and the environmental dimensions, which are less directly involved at the outset, but who may be affected ethically as the implications of the decision unfold.

3   WHAT DOES 'ETHICAL' MEAN?

There are at least five topics within the overall task of being ethical. To be ethical, we must be fair, we must look after our own interests, we must guard the interests of others, we must meet our obligations, and we must be willing to be subject to the same rules as we seek to apply to everyone else. There are other aspects of morality, but they seem less likely to impinge on problems of DSS design.

3.1 The dimension of fairness

The dimension of fairness has been at the forefront of ethical writing since the earliest days. In the Nicomachean Ethics (Aristotle, Ross ed, 1980) the great man, writing about 350BC, devotes half of book five to the topic of justice, and it is clear that he holds justice, lawfulness, and fairness to be close relations. The task of a judge, in cases of rectification, is often one of restoring to the aggrieved party a 'fair' share of whatever was in dispute. Indeed, the word for judge and the word for 'bisect' are nearly the same; that is the task the judge performs. A hundred and fifty years earlier, Confucius (Fung, 1948) had addressed the same kinds of problem, and had arrived at a very similar answer.

There have been contributions on fairness from a large number of writers since then, but we will 'fast forward' to the present. John Rawls ideas seem very mechanistic, in general, but he does offer a useful definition of a fair system (Rawls, Bedau ed, 1969) when he states that it ought to be based on a pair of rules under which (1) each person has the most extensive liberty compatible with the same liberty for everyone. This seems very reminiscent of Kant's universal principle of justice. Rawls second rule (2) requires that social and economic inequalities be arranged to be for everyone's advantage, and attached to positions everyone could aspire to. Midgley has complained (Midgley, Eliott ed, 1966) that Rawls' ideas are insufficiently humane, by simply omitting anything resembling compassion from his writings. This criticism seems valid; the Rawls items I have read are seriously too legalistic. At the same time, his idea of granting rights of aspiration to, and access to the routes toward, prosperity is a good example of what the word fair means in its financial manifestation. Rawls might also argue that compassion is inherently unjust. If we give someone something out of compassion, they have no right to it (presumably), so justice is not being done. Love or brotherhood, perhaps, but not justice.

On the intensely practical side, for many years now, company auditors have reported that an account showed a 'true and fair' view of the firm's financial condition. Until it fell into the hands of the lawyers, and was reduced to being the summation of an assortment of pronouncements by very variable judges, 'true and fair' was a very clear concept. In addition, it was clear that it was not a precise concept. The profit reported was one of a range of numbers within which range all competent evaluators would agree the true profit lay. If the firm tried to report a profit outside that range, the report would be qualified. The principle was based on prudent judgement and fairness to all parties, and, as Kavka points out 'Rules of conduct grounded in prudence and reciprocal in form, connect with two of the most ubiquitous

and reliable of human motivations, rational self-interest and a sense of reciprocity or fair play'. (Kavka, 1986, P310)

In discussions with people in the business communities, Barry (1979) established that the concept of fairness was the central essence of what they thought business ethics was all about. When pinned down, the business people tended to formulate their interpretation of fairness in business in terms very like the silver or golden rule.

In building a DSS, it is surely a vital element to look ahead, to see what conflicts might arise, to consider whether we might be deemed, by public opinion, to have dealt unfairly with people. There are situations in which this is a definite risk, and it is then essential that we should design our DSS so that it will flag a warning of a possible future difficulty, whether of public relations, or employee relations, or even of law.

3.2 The dimension of self

There is considerable agreement that everyone has an obligation to look after himself. This is different from being selfish, a word which is never used in English without a negative connotation which implies that you have failed to contribute your fair share to some common activity. The self is a very important creature to the majority of ethicists, and is to be looked after with considerable care and thoughtfulness. The idea of self-love has been in use at least since Aristotle, but the name seems to have been introduced by Bishop Butler, and comes in a range of available temperatures.

The Bishop (Butler, Darwall ed, 1983) proposed a 'cool and settled selfishness' from which interested actions could proceed, and contrasted this with a 'passionate or sensual selfishness' which resembled a desire amounting to lust for power, honour, or the good or ill of another person. Butler points out that the people need to have love for themselves and also to have love for others. He uses the word benevolence for that, and we will consider it further below. A man will follow courses which will do him good directly in himself, and also that will do good as a citizen, and by the first he will benefit the public by being a more useful contributor and by the second he will benefit himself by increasing society's regard for him. We need both, in his lordship's view. We have an obligation to think ahead and to behave with reasonable caution, because God gave us (but not the other animals) a brain to do that with and we deserve censure if we fail to make good use of it {Sermon 2. Para 10}. God also supplied a conscience in each person which governs and over-rules (if need be) all the other appetites. This is perhaps most explicitly set forth in another book, the Dissertation on the Nature of Virtue (Butler, Carlsson ed, 1964), where the Bishop defines prudence as a 'due concern about our own interest or happiness, and reasonable endeavour to secure and to promote it.'

Thomas Hobbes was active before Butler, and was rather more of an enthusiast for self-love. If Butler favoured a cool version, Hobbes' was too hot to touch. 'The right of nature, which writers commonly call jus naturale, is the liberty each man hath, to use his own power, as he will himself, for the preservation of his own nature; that is to say of his own life; and consequently, of doing anything, which in his own judgement, and reason, he shall conceive to be the aptest means thereunto.' (Hobbes, Sneath ed, 1898, Of Man, Chapter 13).

A third version of self-love is offered by Immanuel Kant, who serves it up on ice. According to Sullivan (1994), he instructs us to regard ourselves as a valuable human being, with capabilities to be developed, just like everyone else. We should not treat other people any better than we treat ourselves. We have a strong moral obligation to build up our own capabilities. This is what he means by treating our abilities with dignity, which means we have to work hard to enhance and to develop them as far as they can be made to go.

The categorical imperative, which is the fundamental ethical instruction Kant issued, is provided in three forms. The second of these, which is the one favoured by Pope John Paul II (La Barge, 1990), states that we have a moral obligation to act so as to treat humanity, whether in the form of our own person or of any other person, always as an end, not as a means. (Sullivan, 1994). The instruction appears to instruct us to treat everyone with respect, which may seem uncontroversial, but also to treat everyone equally, which could easily be interpreted in a misleading way.

The concept of self-love is therefore considered carefully by a number of important writers in the ethical field, and it is clear that we are expected to pay attention to our own welfare as well as to that of others. It must be admitted, as a matter of practice, that the problem of under emphasising our own interests has not been a very major problem in DSS design problems I have been involved with, or have observed over the years. All too often, the idea of self-love, in some kind of corporate format, has been the only explicit goal.

3.3 The dimension of unselfishness or humaneness

Every document calling itself ethical has a prescription in it somewhere that encourages the reader to think of other people besides himself. There have been a range of texts, and a wide variety of recommendations on the importance of this dimension of ethicality. It is the aspect of ethics which is the most controversial in business situations.

The controversy is not confined to business situations. There is a very broad range of opinion among philosophers on this matter. At one end of the spectrum, there are the universal love advocates, led by Mo Ti (Jochim 1980) and Jesus Christ, both of whom would have us love our enemies as well as our friends. At the other end, there are the earthier philosophers like Hobbes, Machiavelli (1993), and Mandeville, who take the whole idea of self interest very seriously indeed. As tends to be true with most spectra, there is a very large group in between, who offer an almost infinite variety of options.

It is contrary to the nature of a business negotiation to adopt the universal love principle. The executive who tried to do unto others as he would like them to do unto him will wind up giving away his shareholders' entire company.

The views of Hobbes and Mandeville are much more likely to enable a business negotiation to occur. Their views involve sympathy to those in need or in distress, but a powerfully self-orientated approach in other cases.

Bernard de Mandeville alleges that all the apparently virtuous actions we take are actually prudential or selfish actions in disguise. In a rather marvellous doggerel saga-poem accompanied by some hilarious but occasionally scurrilous notes, (Mandeville, 1989), suggests that the unselfish actions of people are nothing more than cunningly manipulated acts of self-interest. The manipulators are 'dextrous and skilful politicians', who manage to turn

the vices of the citizens to the public benefit. To my view, the important bit is that he says a person should get no credit for virtuous action, because he has only engaged in that action because of some selfish motive. The man who gives to charity in public, the soldier who is brave from fear of shame, and others are really the victims of a vast conspiracy by the powers that be to achieve the goals of unselfishness from a totally selfish company.

Hobbes had a rather minimalist view of unselfishness, which Kavka (1986, p347) called the 'Copper Rule' to emphasise that it was neither Gold nor Silver. 'Do unto others as they do unto you' he suggests. By all means extend the right hand of fellowship to the other side. You must sincerely try to maintain a moral covenant. But if the other lot back out, you are free to hit them with all the ammunition available. You are certainly not under any obligation to continue to abide by a covenant the other side is obviously not going to stick to.

Thomas Hobbes, despite his apparent role as bete noir of ethical philosophy, had a characteristically robust view about what to do to people who refused to help the needy (Hobbes, Leviathan, Ch15). 'A man that by asperity of nature, will strive to retain those things which to himself are superfluous, and to others necessary, and for the stubbornness of his passions, cannot be corrected, is to be left or cast out of society, as cumbersome thereunto. For seeing every man, not only by right, but also by necessity of nature, is supposed to endeavour all he can to obtain that which is necessary for his conservation; he that shall oppose himself against it for things superfluous, is guilty of the war that thereupon is to follow'. He goes on to say that the state is responsible for making sure that those in need are looked after; they should not be left to the charity of individuals.

In building a DSS, and indeed in handling many other kinds of business transactions, it is the comment of Hobbes which should be in our minds. It is just not good enough to harm those who cannot defend themselves by means of a routine business transaction, in which we simply did not think through the problem enough to realise the damage we were doing. It seems reasonable enough to be fairly tough in business-to-business dealings, subject to obligations (see below), but we ought not squash people who cannot fight back. As DSS designers, the adoption of this moral requirement adds a serious additional item to think about to our already well-laden list.

3.4 The dimension of obligation

In philosophy, an obligation is something we ought to do, but may refuse to do. Ross (1970) suggests we may upset a few people, or even quite a lot of people, by failing to meet our obligations, but the choice to refrain from fulfilling them remains totally in our own hands, and failure to meet them need cost us nothing, not even a sleepless night. Individuals acquire obligations from their roles. As father, I ought to take my son to a football match. As a company director, I ought to give clear guidance to my various managers. As a bishop, I ought to encourage my vicars in their ministries. I acquire the obligations by having accepted the status. In the case of a company, obligations are acquired by signing contracts, but also by making informal promises. A company may even acquire obligations by the simple passage of time. If a firm has financed the village brass band for the last sixty years, it has an obligation to go on financing it.

It is not uncommon for people to feel more strongly bound by an obligation than they feel bound by a legal contract. There is a substantial difference in behaviour among nationalities here. The real issue is communitarianism. If I am a part of a community, I will conform to its rules. Moreover, if I fail to conform to the rules of the community, I will expect to be ostracised or penalised in some other way. This is a powerful conformative force.

Every person, and every entity, needs to be part of at least one system each and most of us are part of dozens. It is partly a matter of prudence. Loners get "picked off", while herds and tribes live on, and so do most of their members. It is partly a matter of foresight. A co-operator will get system help in time of need, a whinging isolate will be left to whinge even as he failed to help the rest. Long-term far-sighted egoism entails trust in other system members. Long-term far-sighted egoism may even entail occasional heroic suicidal action, where the "self" being preserved is "man's own nature". Hobbes set that equal to his life, but it could be his progeny. Or her progeny. The best way to look after number one is to persuade others to behave in a long-term far-sighted egoistic fashion, and to do the same yourself. Some will be persuaded on grounds of duty, some on grounds of benevolence, some on grounds of prudence. Use whatever will work to promote system loyalty, obligation, and support. On historical evidence, egoism is the easiest to "sell", and long-term far-sighted egoism is only slightly harder.

As DSS designers, we are at a disadvantage on this topic. The problem is that we do not, in general, know what obligations the company or person we work for has taken on. Sometimes these are phrased very loosely, and this can come across in incredibly loose terminology when the instructions are being drafted for a new DSS. The only way to cope with this is by means of a series of questions to try to ensure that we do not miss anything large.

3.5 The dimension of universality

Immanuel Kant's famous categorical imperative (Kant, Abbot Translation, 1946) required universalisability before a maxim was allowed. We are prevented from using a maxim unless we would be prepared to require everyone else to stick to it too. Universalisability is a reasonably straightforward, rather stringent, screening system for weeding out policies which fail in some way to support the mission of not making the inequality situation worse.

4   EXAMPLE (A) A DSS FOR FACTORY RELOCATION:-
EXPANDING OUR HORIZONS THROUGH SPACE

We have now looked at the essential ingredients of ethicality, and it is time to try to show how this new dimension can be employed in the typical DSS situation. We will look at a capital expenditure model for factory location in this section, and then at a capital expenditure model for natural resource management in part five. It will be shown that we must expand our horizons. In part four the horizon expansion is through space as we consider the impact of the change on communities in many places. In part four the horizon expansion is through time, as we consider the impact of the natural resource exploitation activity on future generations.

The DSS models we use in handling capital expenditures are very well defined. Unfortunately, because they are highly quantitative in character, they usually miss out some very important parts of the computation. We know quite well how to carry out the net present value calculations, and how to determine the cost of capital and how to estimate the cash flows. There are a few minor controversies on these issues, but by and large the finance community considers the problem pretty completely solved. It is not a contention of this paper that the finance community does these sums wrongly in the mechanical sense. Instead, it is my contention that they forget several important elements of the computation completely.

It is a common occurrence for a multinational company nowadays to consider moving a factory from one country to another, or perhaps moving the production of a product group, so that an old factory is no longer required. When we are working out whether this would be a good move or not, the usual computations show certain labour costs as being discontinued (usually involving a redundancy payment), certain other labour costs as commencing, and various other costs being incurred for the physical movement of various tools from one site to the other. The normal computation would also show that the material costs, being drawn now from a different source, have changed, and that the costs of transportation of the final product have changed. Obviously, the expectation is that the proposed move will only take place if the net present value of the differential cash flows of the project after tax are positive when discounted at the company's after-tax cost of capital. This process is well understood.

The problem is the omissions. I have had an opportunity to examine quite a number of these DCF computations in several companies (45 in 8). The cash flows as stated above are shown, but the cash flows associated with the true costs of winding down the old plant in a humane and ethical fashion are definitely not estimated well. The only recognition that the employees are being laid off is the budgeted figure for a legally required redundancy payment. In some cases, even this is omitted. There is, typically, nothing in the sum which talks about the costs associated with the outplacement service. There is not likely to be any mention of the costs incurred by the company in helping the community to attract new employers. The costs of assigning a good, and rather expensive, manager to the community to help it get over the shock is not included in the DCF sum either.

The reader may say that it is quite in order, mathematically speaking, for these items to be missed out. If the company was not planning to incur these costs, and was really planning to head for the foreign site without taking any of these steps to alleviate the problems caused by the move, the sums they are doing, as described in the last paragraph, are a true and fair reflection of what is going to happen. That point is certainly accepted. We may not like them much, but at least they are doing their sums aright.

My ethical complaint is different. What is usually happening is that the multinational actually goes to quite a bit of trouble on these matters. A good down-sizing manager is sent in to help out. The local outplacement industry is mobilised at the company's cost. Extensive costs are incurred at the site being left which were not included in the capital budgeting analysis. Why do I complain about this, then? Surely if the company is being as helpful as this, there is nothing to get excited about, and we should all be rather grateful that the company is doing its best to help?

The problem is that the omission of these costs from the budget on which the decision was based biases the decision in favour of moving. The costs of moving out of the old site are systematically underestimated. In some instances, the costs of opening up the new site are

even more seriously underestimated. I took the trouble to work it out in one instance, the only instance where I could get any figures for these items. The result in that particular instance was to reduce the decision to move to a net present value of zero from a net present value of (about) a quarter of the cost of the machinery being moved across. A zero NPV still indicates that the action should be undertaken, but it would be very hard to justify the move in view of the management time and trouble that had to be expended to obtain such a minimal reward.

Do I then have any ethical complaint on this topic? I believe so. Case (1) is where the company does not carry out these remedial tasks, and does not budget for them, in which instance it is getting its sums right but is failing seriously in the morality of its actions, as discussed earlier when we looked at the moral injunctions, for instance, of Kant (1954). The company may be said to failed on humaneness and on meeting its moral obligations. It is difficult to tell whether it can be criticised for unfairness; this rather depends on employment alternatives in the community they are leaving. Case (2) is where the company does carry these tasks out, but does not budget for them, and therefore exposes itself to the risk that it has made the wrong decision. This is less likely, but is an even more agonising situation. The firm in this case has failed to look after the dimension of self-love, and may have violated its obligations to its own shareholders. Only in Case (3), when the company takes the actions after budgeting for them, is the situation in ethical balance.

My specific request is that the DSS models which are so commonly deployed for capital investment decision support, especially for evaluating situations in which employment levels will be changed, perhaps simultaneously in several different points in space, should be augmented to take account of the ethical realities. A company which takes its decisions on one set of criteria and actually behaves according to another set is taking serious ethical risks, and also serious prudential risks.

5   EXAMPLE (B) A NATURAL RESOURCE MANAGEMENT PROBLEM:- EXPANDING OUR HORIZONS THROUGH TIME

The second DSS example we will look at is another capital expenditure analyser, which carries out the same computation as the one we have just looked at, but which is normally written to a different design to deal with the special situation of a natural resource company such as a forest products group or a mining concern.

In this situation, the problem of disrupting the lives of people and indeed of entire communities is commonly present in mining projects, but we will not consider this again as it has been covered well enough in the discussion of section four. Instead, we will consider the need to cope with the demands of morality by expanding our planning horizons through time. Far too often, in this kind of situation, the DCF sum is carried out over a very short planning horizon, and this can have an appalling effect on the environments, both physical and human.

It would be helpful to consider an example of this problem. The general corporate tendency, and even more generally the governmental tendency, is to assume that the period of time over which a project should be assessed is the length of time the machine will last. If we are going to buy a computer to perform a particular task, and the machine will be useable for six years, then the cash flow computations are done over the six year period. Occasionally, the

decision may be taken to assume that we will get another machine after the first one fails, in which case the decision is computed over twelve years, but that is done so seldom that it could almost be ignored. If the project is going to generate sales revenues which increase steadily over time from a low base to a stable state in year ten, which is not an uncommon pattern, the futility of stopping the analysis in year six becomes obvious.

Table 1 - Illustrative figures for Example B

Year	Machine	Net Sales 14%	Net Cash	DCF Factor 12%	DCF	Total DCF to date
0	-7500		-7500	1.00	-7500	
1		1000	1000	0.89	893	
2		1140	1140	0.80	909	
3		1300	1300	0.71	925	
4		1482	1482	0.64	942	
5		1689	1689	0.57	958	
6		1925	1925	0.51	975	-1898
7	-7500	2195	-5305	0.45	-2400	
8		2502	2502	0.40	1011	
9		2853	2853	0.36	1029	
10		3252	3252	0.32	1047	-1211
11		3252	3252	0.29	935	
12		3252	3252	0.26	835	558

In this table, we can see that the total net present value to the end of year six is -1898, which suggests that we should not engage in the project. By the end of year twelve, the net present value has gone positive, which suggests that we should engage in the project. Obviously, I have concocted these numbers to make them come out that way, but the problem is very real. Time and again the procedures manuals and the decision support models in use for capital projects have been based on the wrong time horizon, and I have to report that in 35 years of looking at this kind of issue I have not yet seen a British company or an American company use a time horizon that was too long.

In the case of decision support systems, the usual error is that the model is constructed with a fixed format, so that everything has to stop at an arbitrary date, which is often at the end of year ten. In the example above, the project would still have been looking negative (npv =-1211) at that stage. This DSS defect is, however, not my main concern. It is an error, and a serious one, in terms of evaluating the economical situation soundly. In this paper we are concerned about the ethical situation to a greater extent.

The ethical problem is very nearly the same in appearance, in that it depends on the use by the DSS of too short a time horizon. The essence of the problem is different, though, because a large proportion of the ethically important items get missed out altogether because of this.

The morality issue which usually arises in a mining situation is morality through time. If we sink a shaft, drive tunnels hither and yon, and extract non-renewable natural resources then

we are obviously making a choice which will affect all future generations. The oil, or the coal which we take out will not be available to them. If we are mining for gold or diamonds, the assets are more permanent, and the assumption may fairly be made that no restitution needs to be made for the actual product. In both instances, however, the territory surrounding the mine may be permanently damaged, and we have an obligation to repair that or to render the mine safe in some other way. There are parts of the tin mining areas of Cornwall which have had to be permanently sealed off because a millennium of digging has left the land in a completely lethal state.

If the natural resource company is dealing with renewable resources, the situation is potentially less devastating, but we still need to handle the DSS design problem very carefully. If a timber stand is felled, and the timber is put to sound economic uses, the territory can be either left to regenerate or it can be actively replanted. There is controversy as to how this should be done. Active replanting immediately after a felling has been criticised as likely to damage the land for a very long time. Leaving the land for a couple of years and then replanting was the favoured option when last I had the change to discuss the matter with a Forestry specialist.

The consequence of this, of course, is that the DCF calculation must be done with a very long time horizon indeed. The entire production cycle varies from one species to another, but let us suppose a stand matures in fifteen years. The planting takes a year, growing takes fifteen, felling takes one, two years fallow, and restart. A grand total of nineteen years. There is no point in doing the DCF computation for any smaller period than that. If we fail to take these reconstitutive charges into the account, the chances are that someone will simply try to walk away at the end of the felling, leaving the place to do its own regeneration. It is little short of miraculous how well nature has hitherto managed to cope with this behaviour, but there are limits.

If we are going to handle the situation ethically, we need to recall the need to be fair, and to meet our obligations. In this instance, the obligations are to the next generation, or the one after that. As a general rule, in renewable natural resource DSS design, it would make sense to establish a regulation that the time horizon must equal the growth interval plus five years.

In the case of non-renewable natural resources, the time horizon is extremely hard to predict. The mining company will be hoping the mine will last forever, and most people will completely agree with them. The only way to handle the need to build a restitution fund in this situation is by means of a levy on the cost of the mine, which is used to build a capital sum which can be used by the local government to put the land back to its original format after the mine stops. If the mine is a failure, there may not be enough in the kitty to restore the land, but there will at least be something to alleviate the situation.

Although this section has been written in the context of mining companies, there is reason to believe that DSS designers working for routine manufacturing or service businesses should think in terms of long horizons as well. It is only necessary to look around the derelict buildings which clutter every city to appreciate that the termination sequence of a plant, or project, or company was not taken into consideration.

There are very strong practical reasons why this was so. The man who started the business had no intention of ever stopping. The company, in his dream, would go on for ever, and would eventually be a mega-business managed by his great-grand children. Making provision for cleaning up the mess after it folded would never enter his head. Also, when, at some future

date, the business got into difficulties, and when it became increasingly obvious that there was a serious risk of collapse, the capacity of the firm to make an provision for cleaning up would be very seriously impaired. In the last few years before a corporate collapse the managers usually engage in all sorts of rather desperate manouevres to keep the show on the road, each of which consumes a portion of the ever-shrinking resource base. Closing the place down while there is still enough to give everyone a decent payout happens very rarely indeed. Accordingly, it is suggested that there is a need for the local authority or some similar government agency to take a role, building a fund for each business out of revenues raised from it, which will enable the restoration of the status quo ante.

From the point of view of our main concern, the ethical aspects of DSS design, the issue is one of fairness from one generation to the next, and of meeting obligations. We may not have a legal obligation to restore the land on which our factory has been sitting for fifty years, but there is a very strong case to be made for a moral obligation. It is a contention of this paper that DSS designers have a professional obligation to consider, in each instance, whether issues like this are applicable to the DSS they are working on.

6 CONCLUSIONS AND ACTION RECOMMENDATIONS

In moving toward the formulation of my conclusions I would like to introduce two more powerful authorities to set the scene. They are Adam Smith (1969) and W Ross Ashby (1957). Everyone knows about Adam Smith, as the inventor of economics and especially of the "invisible hand" which makes markets operate. However, my reason for citing him is the other major work he wrote, which makes clear that the invisible hand cannot work unless the society in which it is placed is stable and humane. Ross Ashby contributed the separate but valuable insight that a system cannot be controlled except by another system possessed of the requisite variety, which normally means the control system will be as complex as the original one.

Adam Smith wrote two books which should be read together. Too often the only work of his that is quoted is the single paragraph about the invisible hand which guides the butcher and others to serve up dinner. In fact his two main books were his lecture notes for two of the four segments of the same university course, and ought to be taken as a totality.

Smith's Wealth of Nations is one of the most famous texts ever written. His Theory of Moral Sentiments (Smith 1969) is less well known, probably even less often read, but it has a very powerful message too. You really cannot understand the picture in the Wealth of Nations unless you place it in the frame of the Moral Sentiments. This is a long book, because it explores, most elegantly, the byways of all the meanings which the virtuous words and the vicious words can take on, and explains why one is more or less meritorious than another. The role of conscience, the authority of the general laws of morality (3.5), and the importance of self-control (4.2) are all discussed at sufficient length that even the slowest of his undergraduate students would probably have got the general idea.

The main point is that he was setting forth a very coherent and detailed set of assumptions about the kinds of behaviour that he expected people to engage in. The invisible hand which appears in the Wealth of Nations to serve up dinner, appeared also (fifteen years earlier) in
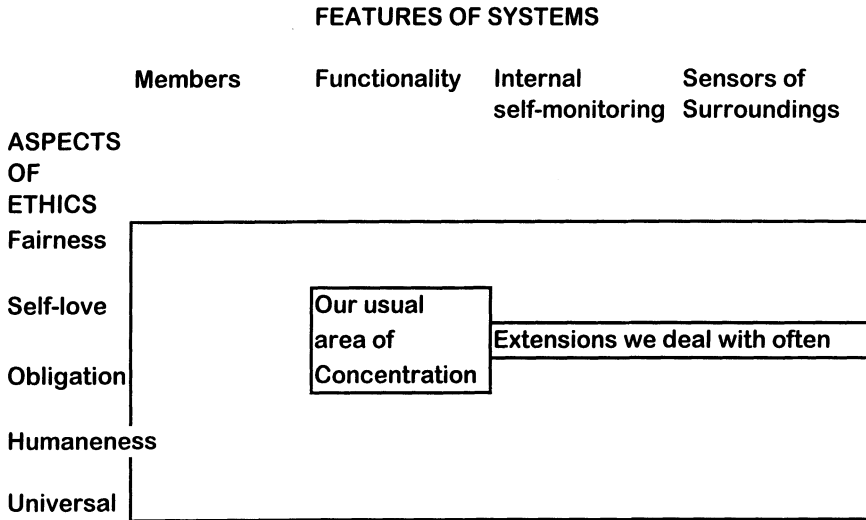
part 4 of Moral Sentiments, this time in the role of sharing the landlord's crops out equally amongst the population. Those criticising Smith's writings seem to be founding their complaints on a single paragraph, and ignoring the very specific, interpersonally ethical, relationships which he consistently promoted, and within which setting the business operations in Wealth of Nations were to be assumed to be set.

It seems to me that we have a comparable situation now, a quarter of a millennium later. The technically elegant scheme (Smith's economics) could only work reliably in the defined social context (described in Smith's book on morality). In our own day, we have to set our technically elegant DSS designs and plans in as sound a framework of the total social context as can be arranged, or they may well give us strange, and perhaps shocking, results.

The second topic is on the dimension of control. A very important work in cybernetics was the pair of books by William Ross Ashby on the design specifications for a brain. In this he drew up an effective proof of a most significant theory of cybernetic control which is a rather complex concept, but which has a very simple result. He called it the 'Law of Requisite Variety'. If you have a complex system with N dimensions to its complexity, you cannot control its behaviour unless you have a control system which also has N dimensions to its complexity. As a simple example, consider a car. It can travel at a speed from (let us say) -30 to 110 miles per hour, where a minus speed entails reversing. It can be aimed, at any given moment, a few degrees to the left or to the right of its existing course. There are two dimensions of complexity, and you therefore need two dimensions of control. Obviously, these are the accelerator/brake/gearchange system and the steering wheel respectively.

When you consider more complex systems, the problem of creating a control system becomes much more serious. A system which resembles a company has a very large number of dimensions of control. Attempts to list them have never been very successful, but they would include product line range, production methods, region serviced, region from which supplies are drawn, people employed, and customer groups sought. Each of these may be subdivided in certain instances where this becomes important. The task of creating a control system is made more complex because the firm has to ensure that there is a control system for each of these dimensions, and also that there is a control system which achieves co-ordination among them. The task is probably solvable in most cases at the level of a company, though there are many instances where the task has not actually been solved at that level. When we try to add the attainment of ethical standards to the list of dimensions, there will be need for a further set of control measures to help the system understand whether it is meeting its own standards or not. Very few company level systems have even tried to deal with this dimension so far.

With that introduction to the conclusion, let us move to its statement. It is shown in figure 1. Across the top are to be found the features which systems, including DSS, all possess or ought to possess. Down the side are aspects of ethics, which are dimensions I am contending we, as DSS designers, ought to be thinking about in some situations, and certainly in a lot more situations than we actually do think about them. Diagrams are always somewhat blunt instruments, but the main point is that the area of our normal concentration of effort in DSS design work is the small, tadpole-shaped area in the middle.

FEATURES OF SYSTEMS

	Members	Functionality	Internal self-monitoring	Sensors of Surroundings

**ASPECTS
OF
ETHICS**

Fairness

Self-love

Obligation

Humaneness

Universal

Our usual area of Concentration	Extensions we deal with often

Figure 1 - The Features of Systems and The Ethical Aspects which we typically handle well, and which we typically handle badly.

I think we can reasonably claim to do a good job of handling the functionality of our models, and that we are doing a good job of coping with the ethical requirement to be prudent and look after our own interests. We also do a good job, if the corporate bosses remember to tell us about them, of dealing with the obligations which our company may have to other firms. We also, I am suggesting, do a good or at least fair job of handling self-love and obligation as ethical aspects of our design work of the monitoring systems we build into DSS for their internal cybernetic management, and we spend a lot of effort on data-gathering and on extracting knowledge from experts to enhance our designs.

There is a large territory in the diagram which we do not often deal with. I am prepared to believe that, if we worked at it for a while, we would be able to find instances of DSS which dealt with almost every box on the chart. At the same time, I feel secure in claiming that much of this territory is severely under-explored. The column about membership is conspicuously empty. We deal with membership from a rather narrow viewpoint, counting ourselves, our fellow employees, our company's customers, and perhaps our company's suppliers as the universe of discourse. What about the people we do not or will not sell to? What about the people we do not or will not hire? Sometimes, I concede readily, these are irrelevancies, but I hope I have shown earlier that they are not always.

The other areas in the diagram which are under-explored are the rows labelled by the ethical terms "fairness", "humaneness", and "universality". Taking the first two together, I am suggesting that we just do not have a way of handling these dimensions at the moment. It

would be theoretically possible to write a DSS which exhibited justice if you could tell it what was just in its particular milieu. It might even be possible to write a DSS to support the same decision which exhibited humaneness. However, a DSS that tried to do both would probably develop symptoms which could be anthropomorphised as executive stress. To be humane, it would have to give at least one person more than their entitlement, which is unjust (Rawls 1991). It is not suggested here that the handling of these dimensions of ethics needs to be included within the DSS itself, but that these dimensions of morality should be considered during the design process in order to meet the point made by Adam Smith above, that there is a context within which the system will work, and it probably will not work in a different one.

The last point is universality. Immanuel Kant instructs us to adopt rules of behaviour only if we would be satisfied that they should apply to everybody. At the present time, it would be only very slightly unfair to suggest that our maxim for the production and design of decision support systems is that "we will undertake to make systems which work well in functional terms, which can be adapted to new circumstances, which incorporate sensitive sensors about environmental change, and which are orientated towards our corporate profits and meeting our obligations, especially legal ones, and which do not take any account of what this does to the rest of the world's population of people, plants, or animals or their successors into the future". Is that really what we mean? Is it really what we want to mean? I am contending that it is what we are actually doing, by default.

7 REFERENCES

Aristotle (1980) *The Nicomachean Ethics.* Oxford UP, Oxford.

Ashby, W.R. (1957) *An Introduction to Cybernetics.* Chapman & Hall, London.

Barry, V.E. (1979) *Moral Issues in Business.* Wadsworth, Belmont, California.

Beer, S. (1979) *The Heart of Enterprise.* Wiley, New York.

Butler, J. (1983) *Five Sermons preached at the Rolls Chapel and a Dissertation upon the Nature of Virtue,* {ed. Darwall, S.L.}. Hackett Co, Indianapolis, USA.

Carlsson, P.A. (1964) *Butler's Ethics.* Mouton and Co, Hague.

Fung, Y. (1948) *A Short History of Chinese Philosophy.* Macmillan, New York.

Hobbes, T. (1898) *The Ethics of Hobbes, as contained in selections from his Works.* Ginn & Co, Boston, USA.

Hutcheson, F. (1971) *Illustrations on the Moral Sense.* Mass Belknap Press, Cambridge.

Jochim, C. (1980) Ethical Analysis of an Ancient Debate:- Moists versus Confucians. *Journal of Religious Ethics* **8** (1): pp 135-147. Univ of Tennessee, Knoxville, TE, USA.

Kant, I. (1954) *Kant's Ethical Theory.* Oxford, London.

Kant, I. {Abbot, T.K., translator} (1946) *Fundamental Principles of the Metaphysic of Ethics.* 10th edn. Longmans Green, London.

Kavka, G.S. (1986) *Hobbesian Moral and Political Theory.* Princeton UP, Princeton, NJ, USA.

La Barge, J. (1990) Economic Systems and the Sacramental Imagination. In: Gower, J.F. (Ed.) *Religion and Economic Ethics*, pp. 151-172. University Press of America, Lanham, Maryland, USA.

Machiavelli, N. (1993) *The Prince.* Wordwort Reference, Ware, Hertfordshire, UK.
de Mandeville, B. (1989) *The Fable of the Bees.* Penguin Books, London.
Midgley, M. (1966) Duties concerning Islands. In: Elliot, R. (Ed.) *Environmental Philosophy*, pp. 171-185. Univ Queensland Press, St Lucia, Queensland, AU.
Rawls, J. (1969) The Justification of Civil Disobedience. In: Bedau, H.A. (Ed.) *Civil Disobedience*, pp. 240-255. Pegasus, New York.
Rawls, J. (1991) Definition and Justification of Civil Disobedience. In: Bedau, H.A. (Ed.) *Civil Disobedience in Focus*, pp. 103-121. Routledge, London.
Ross, R. (1970) *Obligation: A Social Theory.* Univ Michigan, Ann Arbor, MI, USA.
Smith, A. (1969) *The Theory of Moral Sentiments.* Arlington House, New Rochelle, NY.
Sullivan, R.J. (1994) *An Introduction to Kant's Ethics.* Cambridge Univ Press, Cambridge UK.

8   BIOGRAPHY

Andrew McCosh has been a member of IFIP since 1984. He was professor of management accounting at Manchester Business School from 1971-1985, and Dean of the Faculty for part of that time. He was professor of finance at Edinburgh University from 1986-1995, and Head of Business Finance, Accountancy and Economics for part of that time. He has published forty papers and seven books on systems and on finance.