December 2003

On Ethics and Decision Support Systems Development

Rob Meredith
Monash University

David Arnott
Monash University

On Ethics and Decision Support Systems Development

Rob Meredith[a] and David Arnott[b]

Decision Support Systems Laboratory, School of Information Management & Systems, Monash University
PO Box 197, Caulfield East, Victoria, Australia, 3145
[a] Rob.Meredith@infotech.monash.edu.au
[b] David.Arnott@infotech.monash.edu.au

Abstract

The ethical aspect of decision support systems (DSS) is an important area of concern for developers and users alike. Such systems impose frameworks and structures upon the cognitive decision making process to a greater or lesser extent, requiring the developer to anticipate, if consideration is given at all, the ethical questions that the decision maker might face. However, the level of research in DSS ethics is disturbingly low. We turn to the area of medical decision support where the four bio-ethical principles of beneficence, non-maleficence, autonomy and justice have been identified as a useful framework for ethical medical DSS. We believe this framework is useful for DSS in general, and present a call to arms for further research into DSS ethics.

Keywords

Decision Support Systems, Ethics, Autonomy, Beneficence, Non-Maleficence, Justice

Primum non nocere
(Above all, do no harm)
Galen, c. 130CE – c. 201CE

Introduction

When decision makers are faced with a decision situation, they often have to contend with a number of competing factors to make a 'good' decision. Not the least of these is whether the outcome of the decision process is in accord with not only their values and principles, but whether they fit into the broader values and principles of other stakeholders and society at large. Ethical decisions are something that most of us strive for.

Decision support systems (DSS) is the area of the information systems discipline that is devoted to supporting and improving managerial decision-making. Over time the majority of DSS research has focused on the application of new technology to managerial tasks at the operational and tactical management levels (Eom & Lee, 1990; Mallach, 2000; Raghavan & Chand, 1988). In terms of contemporary professional practice, DSS includes personal decision support systems, group support systems, executive information systems, online analytical processing systems, data warehousing, and business intelligence (Shim et al., 2002). In this paper we will be focusing specifically on personal decision support, that is, systems designed to aid an individual decision maker (or at least a very small number thereof)

with a single specific decision task, and the ethical issues that developers need to grapple with in undertaking this kind of development.

Given the importance of ethical decision making, consideration of ethical issues related to the tools that support that decision making is also important. Whilst ethics and information technology in general has been discussed for a number of decades now, there is very little material addressing the specific ethical issues related to supporting decision makers with technology. Perhaps not so surprisingly, it is the application of DSS technology to medical decisions that has received the greatest attention from ethicists. Given that the debate in medical DSS is so much more advanced than DSS in general, drawing as it does on the vast body of work in bio-ethics, we believe that it may be useful to adopt the ethical principles governing medical practice as a tool to help us understand the principles of ethical DSS practice.

The Nature of Decision Support Systems

Decision support systems generally address ill-structured decisions that are of considerable importance for the decision-maker and the organisation. It is almost impossible a priori to specify the system requirements in such an environment. In some cases, the requirements problem is alleviated somewhat when the personal DSS is developed by the user themselves. This saves explicitly expressing the tacit task understanding developed by the user during requirements elicitation and system use. In this paper, however, we concentrate on personal DSS where the user and developer are distinct people.

The environment of DSS is subject to significant change and even if the system requirements have been specified with some accuracy at the start of the project they are likely to change significantly over time. This system change can occur because of changes in personnel, organization structure and processes, external competitive pressures, and most importantly, changes in the decision maker's cognitive strategies and task understanding. Systems analysts tend to adopt an evolutionary development strategy to cope with this environment.

Evolutionary development in decision support was first hinted at by Meador and Ness (1974) as part of their description of middle-out design. This was a response to the top-down versus bottom-up methodology debate of the time that concerned the development of transaction processing systems. Courbon et al. (1978) provided the first general statement of evolutionary DSS development. In what they termed an "evolutive approach", development processes are not implemented in a linear or even in a parallel fashion, but in continuous action cycles that involve significant user participation. As each evolutive cycle is completed the system gets closer to its final or stabilised state. Courbon argued that the evolutive cycles should be continuous and as rapid as possible as DSS exists in an environment of continuous change. The functionality of DSS is thought to evolve over a series of development cycles where both the client and the systems analyst are active contributors to the shape, nature and logic of the system. Keen's adaptive design model (Keen, 1980) is the most cited exposition of this cycle. Evolutionary development remains at the core of current DSS textbooks (Mallach, 2000, for example).

Within this evolving systems environment the user of a DSS has considerable autonomy with respect to how they make a decision and what support they receive. This places the DSS systems analyst in a very different relationship with the client than their colleagues who develop operational systems. The ultimate users of operational systems tend to be low in the

organisational hierarchy and have no choice as to whether they use the system or not and have no involvement in the design and development of the system. The technologies available to the DSS have multiplied over the last 20 years (Power & Kaparthi, 1998). In addition to spreadsheet, modelling, and database tools, DSS are constructed using executive information systems (Suvachittanont, Arnott, & O'Donnell, 1994), OLAP tools (Thomsen, 1997), data warehouses (Gray & Watson, 1998), and the World Wide Web (Kimball & Merz, 2000).

The nature of support provided by these different technologies can be characterised according to a continuum of passive through to normative support (Jelassi, Williams, & Fidler, 1987; Keen, 1987), with most systems sitting somewhere in the middle. Passive decision support tends to place the emphasis more on the decision maker to control the decision process, whilst normative support imposes a structure and process on the decision maker regardless of their preferences or normal style of work. Passive support tends to consist of the provision of information to the decision maker, leaving it to him or her to assimilate and manipulate that information to arrive at an appropriate course of action. Normative support adopts decision theoretic principles and enforces an 'ideal' process that in many cases takes the decision maker out of the loop. Knowledge-based systems are classic examples of this approach to support, where the system itself makes a judgement on the best course of action given certain inputs. Most decision support systems, however, tend to sit somewhere in the middle, providing some of the structure and process support of normative systems whilst attempting to retain the important element of the human decision maker in the process, ensuring that learning takes place that enables the decision maker to make better decisions in the future. These systems that offer more structured support whilst still respecting the autonomy of the user to control the process are labelled 'active' decision support by Keen (1987). Active decision support implies a more overt intervention into the decision process, which exposes the developer to a range of ethical issues.

Ethics and Information Technology

As a profession, information technology is still in its infancy, when one compares to, say, law, medicine, architecture, or even engineering or education. A large focus of the discussion on the relationship between ethics and information technology has necessarily been on the role of the developer or practitioner as an ethical professional. Mason (1995), for example, argues that the practitioner needs to carefully consider the decisions that they make in the light of the broader impact upon society. Laudon (1995) argues that technology is as much a part of society (and therefore as ethically laden) as any other social phenomenon. For example, the issue of technology putting people out of work is pointed to as an impact that some developers and system owners fail to consider. This debate was particularly virulent in Australia in the 1980s when the proliferation of automatic teller machines was having an impact on the number of bank employees and branches. The typical approach to the issue of the practitioner as a professional is the establishment of professional bodies and associations, and codes of conduct, and there are a number of theoretical frameworks for tackling how one develops such a code (Walsham, 1996, for example).

However, ethical issues extend beyond due diligence on the part of the IT professional during development. The social impact of technology needs to be considered, and issues such as privacy (including users and third parties), the availability of paedophilic material, copyright protection and advertising on the Internet are examples of the kinds of issues society at large needs to consider – not just practitioners (Berleur, Duquenoy, & Whitehouse, 1999; Conger

& Loch, 1995). In these cases, the moral responsibility for the social impact of technology extends beyond just the practitioner: when considering the use and purpose of technology, the responsibility resides with system owners and clients, and in some cases, users as well. This is particularly so when technology is immorally or illegally utilised by users (Banerjee, Cronan, & Jones, 1998).

Whilst a large part of the moral responsibility for the use of technology belongs with system owners and users, it often falls to the practitioner, as a professional, to highlight potential ethical issues in proposed systems. Unfortunately, many IT professionals lack the communicative skills, and the ethical training to be able to engage in an ethical dialogue (Conger & Loch, 1995).

The locus of responsibility is a fundamental issue for ethics and technology. Indeed, the idea of moral responsibility sits at the cornerstone of any ethical debate. Without responsibility for an action residing with a person, then we cannot label the action (and its outcomes) as good, bad or otherwise. They are 'happenings' or accidents rather than moral acts. We see this principle at work in the legal system, where it is incumbent upon a plaintiff or prosecutor to demonstrate intent or *mens rea* on the part of the defendant.

With technological ethical issues, we can see that it can be quite difficult to apportion blame or responsibility when there are so many different actors and stakeholders. The determination of that locus will differ from project to project, system to system, issue by issue. It seems apparent, however, that it is beholden on the IT practitioner, as an expert professional, to ensure that such issues are explored prior to, rather than during or after, an ethical dilemma, and that the relevant actors and decision makers are aware of their responsibilities.

This is more important for some kinds of systems than others. Whilst bank automatic teller machines, or supermarket bar-code scanners, can have a social impact in that they may put people out of work, their use on a day-to-day basis tends not to raise any other particular ethical issues. However, where systems are designed to undertake actions autonomously of their developers, owners and users, or where a system contributes significantly to a decision made by a user, then a large range of potential ethical dilemmas might arise. This is of particular interest to computer scientists interested in the field of artificial intelligence. Lucas (2001) and Dowling (2001) both point out that Asimov was one of the first to codify a set of rules for autonomous systems, albeit fictionally, with his laws of robotics.

Responsibility implies autonomy and free will. The corollary of this is that autonomy and free will carry with them moral and ethical responsibility. If an artificially intelligent system makes a decision, and causes an action to result, who bears the moral and ethical responsibility? The programmers? The system owners? No artificial intelligence system yet has the actual intelligence to comprehend ethical and moral dilemmas and make appropriate decisions, but given that the programmers and/or owners have ceded some control and development of the system to the system itself, it can be difficult to lay the responsibility solely at their feet. The concepts of autonomy, trust and responsibility become more problematic as the system is more active in the decision making process (Dowling, 2001).

Ethics and Decision Support Systems

Whilst decision support systems do not, generally, adopt the same level of autonomy of systems discussed by AI philosophers, they do, to a greater or lesser extent, usurp or impose structures upon the autonomy of a human decision maker. The ethical issues faced by decision support systems, therefore, are a super-set of the issues for non-autonomous

information technology. Given the popularity of data warehouse, business intelligence and other decision support systems, it is unfortunate that the ethics of decision support as a specific topic has received very little attention in comparison to the issues of privacy and other general IT ethics issues.

Indeed, there is a major gap in the literature on this topic. A search of the journal *Decision Support Systems* on the Science Direct website (http://www.sciencedirect.com) for the keyword 'ethics' in the abstract, title or key word list of any article since the January 1995 edition (vol.13, no.1, the earliest edition available on the site) yielded zero results. A more extensive search of the entire text of each article from the same period only yielded ten papers, none of which addressed the topic of ethics directly. A search for the same term in either the citations or abstracts of articles in *Decision Sciences* on the Proquest website (http://www.bellhowell.infolearning.com/proquest) yielded just one result: a paper published *in 1981*. This paucity of published research and debate on the ethics of decision support in two of the discipline's premier journals is disappointing.

This doesn't mean that the topic has been ignored totally. Some of the issues raised include the fact that a decision support tool embodies a particular philosophical approach to decision-making - for example, is it ethical to quantify certain values, such as those we place on human life, or how we manage risk (Johnson & Mulvey, 1995)? Johnson & Mulvey also address the issue of the locus of responsibility for outcomes resulting from decisions made based in part on advice provided by a decision support tool. Their answer to the question is that the developer should have similar responsibilities as any other professional or expert who is hired for their advice. That is, developers should bear responsibility for the quality of the advice their systems provide, including raising and establishing standards and norms for the ethical use of their systems.

A related issue to that of the system embodying a particular philosophical approach to decision making, is whether or not the *correct* decision is being supported, raised by Chae, Courtney and Paradice (2002). They point out that not only is the design of a decision support system not value neutral, it is actually "heavily value laden". Since values have an important role to play in determining whether or not a situation should even be considered a problem, let alone what an appropriate solution might be, ignorance of the various stakeholder value positions in a decision problem can, in fact, lead to the wrong problem being supported. Involvement of various stakeholders in the decision is important during DSS development to ensure that this doesn't happen.

The issues discussed by Johnson and Mulvey, and Chae, Courtney and Paradice are relevant regardless of the kind of decision support tool or approach. However, as discussed earlier, the more control over a decision a support tool has, the more relevant the issue of responsibility becomes. Fox (1990), addresses the issue of expert systems, specifically those used for safety-critical decisions such as those in nuclear power plants, or hospitals. In these cases, the system has a significant level of autonomy to make decisions and undertake corresponding actions. If an ethically questionable decision is made, the moral culpability potentially resides with the system itself. To address this dilemma, Fox suggests that all decisions made by a safety-critical expert system should be subject to possible human intervention. That is, the system should be flexible and robust, to deal with as many unforeseen permutations of the decision task as possible, as well as being accountable to a human decision maker. This allows the moral responsibility to reside with an entity that is morally accountable.

Ethics and Medical Decision Support

It is perhaps not surprising that the most significant contribution to ethics and decision support comes from the medical fraternity. For a long time, medical practitioners have utilised a number of different tools and technologies to support either diagnosis or treatment, often relying on other people to operate and interpret the output of these tools. Examples include radiologists, pathologists, physiotherapists, and all manner of specialists, each of whom provide expert advice that contributes either to diagnosis or treatment of a medical problem. The recent increase in the consideration of medical ethics, partly attributable to the patient rights movements and increased malpractice litigation, has meant that the role that others play in the decision making process has been closely considered (see Goodman, 1998, for example). This, in turn, has meant that the ethical aspects of tool use, including medical decision support, has been scrutinised more than it has for decision support systems in general.

With so many potential participants in the typical medical decision-making scenario, the issue of the locus of responsibility is just as problematic, if not more so, than for ethics and information technology. As Snapper (1998) points out, the responsible entity in a scenario where a physician, patient, specialist, medical technician and a medical device are all involved in the decision making process is very difficult to determine. It could be any of the people or companies involved in the development, use, ownership and maintenance of medical devices, as well as the attending physician. When one adds into the mix the fact that the device may be autonomous and be replacing a responsible human professional (Snapper points to the example of heart monitors), the issue of just who is ethically and morally responsible for an action becomes highly problematic.

In the case of computational devices where a decision is either made by the system, or based on the output thereof, Snapper suggests that the solution is to adopt a similar view to the situation where a physician relies on advice from non-computerised sources such as human consultants. In these situations, responsibility is shared amongst the various professionals. This doesn't, however, answer the issue of who exactly is responsible in the case of autonomous computer systems. It seems counter-intuitive that a computer system or any other tool, lacking consciousness, could be held morally responsible for the outcome of a decision that it supports.

One solution is to ensure that final accountability for a decision lies with a conscious, competent, intelligent agent. That is, the autonomy of a human decision maker must be respected (Abbott, 2001; Collste, Shahsavar, & Gill, 1999; Fox, 1993; Miller & Goodman, 1998; Snapper, 1998). That might mean a human supervisor or reviewer in the case of more prescriptive systems such as expert systems (Fox, 1993), or recognising the idea that a decision support tool is designed to enhance the autonomy of, rather than replace, a human decision maker, in the same way that a stethoscope enhances, rather than replaces hearing (Miller & Goodman, 1998). Snapper provides some empirical support for the argument that decision support tools that provide control of the judgement process are preferred by decision makers. In a study of diagnostic support tools, physicians preferred those tools that provided recommendations along with case histories, rather than numerical probabilities of the correctness of a particular diagnosis. Whilst it is possible that the former type provided better 'peace of mind' for the physician in terms of justification of the recommendation, Snapper argues that it allows for at least the impression of greater control over the judgement process.

Miller and Goodman make the assertion even stronger, stating that a medical decision support tool should never replace a human decision maker: "It must be possible for the user to interpret and even override the data generated through the use of … a decision support system." They point out that there are two corollaries of this. The first is that user interface issues are very important since the ability of the user to interpret the output of a decision support tool is directly related to the way in which that output is presented. The second is that users of a decision support tool should be appropriately trained to understand the use and output of the system, just as users of other diagnostic tools such as imaging equipment undergo training in their use and interpretation. They further state that inappropriate use of decision support systems occurs, not only when the users misunderstand the applicability of the system to a particular situation, but also when such systems intrude in a negative way upon the social structures in place that are designed to assist in the decision making process. The autonomy of people in the decision situation is restricted "when we allow socially productive and respectful relationships to be sullied, or their participants to be taken advantage of." In a field where positivism and normative decision processes are dominant, we see recognition of the fact that technology must be subordinate to social, humanist considerations.

In their discussion of the ethical aspects of a DSS for diabetes care, Collste *et al* (1999) highlight autonomy of the decision maker as important, but go further and point to the four principles of bio-ethics described in Beauchamp and Childress (1989): beneficence; non-maleficence; autonomy; and justice. Whilst Collste *et al* were specifically discussing medical decision support, we believe that there are enough parallels with general decision support systems development to argue that these four principles should apply there as well. Certainly medical decision support systems tend to be a sub-class of personal decision support – there are usually only one or two users of such as system, which is targeted towards a specific decision problem. In the section below, we will show that these four bio-ethics principles can be applied, to a greater or lesser extent, to personal DSS at large.

Another reason for our belief that these principles are applicable is that, of all classes of IT professional, the personal DSS analyst comes closest to being a combination of clinician and practitioner. This is because the development process is client and problem centred; the development is oriented to decision pathology and health; development involves charging fees for services; and the ethical/legal responsibility is to avoid malpractice (Schein, 1987, p.68). The personal DSS analyst tends to work with a small group of clients (usually one) and forms closer professional relationships with them than developers of large scale operational systems.

Another distinguishing feature of DSS development and use is the impact of the system upon the cognitive strategies and structures of the user. Whilst an operational system has some impact upon its users in terms of understanding and task approach, the degree to which a DSS has an impact on the user's cognitive strategies and structures is much greater due to the uncertain, unstructured nature of the task. The intervention of a DSS developer or decision analyst into the life of a decision maker is, whilst generally not as strategic as the life and death interventions that a physician might be called upon to perform, similar to the intervention of a physician to a patient. The principles governing ethical practice on the part of physicians are a useful lens, therefore, for understanding the principles that should govern DSS practice.

Four Ethical Principles for DSS

Beneficence and Non-Maleficence

Whilst Beauchamp & Childress argue that these are separate principles, they are essentially complementary concepts. Literally, beneficence means to 'do good', whilst non-maleficence means to 'do no harm'. One is an act of commission, the other one of omission, both aimed at ensuring that, in the medical setting, the 'good' for the patient is maximised. Beauchamp & Childress specify four directives that can be derived from these two principles:

1. One ought not to inflict evil or harm (non-maleficence)
2. One ought to prevent evil or harm (beneficence)
3. One ought to remove evil or harm (beneficence)
4. One ought to do or promote good (beneficence)

Certainly, our aim as decision support systems developers is to assist decision makers as much as possible. Within the context of a decision problem, we would try to see that any intervention on our part has a positive, rather than a negative effect. As professionals, we should avoid introducing negative factors or processes like complexity or exacerbating the cognitive biases of the decision maker (Arnott, O'Donnell, & Grice, 1993). This holds true whether the negative influence is as a result of our own actions (directive one), a result of external contextual factors (directive two), or as a result of some characteristic of the decision maker themselves (directive three). Our aim, as professionals, should be to always "do or promote good", whilst minimising any negative influences on the decision maker.

Undoubtedly, DSS have a role to play in assisting decision makers – this is, after all, their *raison d'etre*. Beauchamp & Childress, however, highlight an interesting point about the nature of beneficent acts. That is, it is a fine line between acting with beneficence, and slipping into paternalism. This has relevance, particularly for the more prescriptive approaches to decision support, including expert and other relatively autonomous systems.

Even if, in our view as developers, we are satisfying the principles of beneficence and non-maleficence by removing complexity, or minimising the effects of cognitive biases or other negative influences, we may in fact be impinging upon the autonomy of the system user as a decision maker. Our actions can have more than one consequence, and we must be as aware as possible of the full consequences of our actions. It is only in the light of this more complete assessment of our actions that the principles described here should be considered. As we will see, the autonomy of DSS users is important. Unfortunately, a paternalistic approach is not unusual when technologists attempt to act in accordance with the principles of beneficence and non-maleficence. As a result, in the long term, these principles can be defeated.

Autonomy

Autonomy is ultimately about respecting the right of an individual or group to self-determination. Not only is this important as a fundamental human right, it also is a pre-requisite for ethical and moral responsibility. The criminal justice system, as mentioned earlier, places a great deal of importance on the difference between someone acting autonomously, and someone whose actions where the result of influences beyond their

control. In the latter case, otherwise criminal acts are considered to be either less severe (eg. manslaughter versus murder), or are dismissed as criminal altogether. In a political and social sense, there are countless examples of the individuals and groups whose autonomy has been impaired, resulting in gross violations of the beneficence and non-maleficence principles. Just as autonomy is important in the political and social arena, it is a fundamental principle that should underlie support for all decisions, medical or otherwise. Just as physicians should uphold the autonomy of their patients to ensure that their rights are respected, decision support systems developers should uphold the autonomy of the users of their systems to decide for themselves their own course of action.

Beauchamp & Childress state that there are three important criteria for an act to be considered autonomous:

1. The act had to be intentional, a result of an exercise of the will, implying competence on the part of the decision maker to make decisions.
2. The act had to be a result of a decision based on informed understanding.
3. The act had to be free of controlling influences.

In other words, autonomy has aspects of competence, where the decision maker has the requisite skills and abilities to make the decision to act; information, including disclosure to, and understanding by, the decision maker such that they have an informed understanding of the situation and the consequences of acting; and consent, in that they voluntarily commit to the action decided upon. Decision support systems directly impact upon all three of these criteria.

It is perhaps easiest to see this for the first two criteria of competence and informed understanding. Where decision makers lack the skill to make a decision, or process information in such a way as to achieve an informed understanding, a decision support tool can help provide the structure to walk the decision maker through the decision process, or augment the information processing abilities of the user to comprehend fully the decision situation.

The second criterion of informed understanding has long been a goal of decision support. Keen argued in 1980 that user learning, where a DSS user gains insight and understanding into and about the decision problem, is an integral part of successful DSS development, as shown in his now famous framework for adaptive DSS development (Keen, 1980). He goes so far to say that if any of the aspects of the framework are missing, including the user learning loop, then the system is not a DSS in the true sense of the term.

The third criterion, that of being free of controlling influences, poses some problems for a DSS, since it is, in itself, an influence upon the decision-maker. Indeed, if the system or tool had no influence, it would not be of any use. Clearly, the third criterion requires some modification, and Beauchamp & Childress acknowledge this by arguing that the third criterion can never really be achieved. They argue that the standard should be an act free of *excessive* controlling influences, that is, a decision maker should be satisfied themselves that they are voluntarily exercising their free will, without the sense that they are being manipulated or forced to do something that they don't wish to do. In enhancing a decision maker's autonomy by assisting them through augmenting their information processing abilities, or guiding them through a decision process, a DSS shouldn't become an excessive, controlling influence. In other words, in an effort to boost the autonomy of the user, the support provided doesn't descend into paternalism, thereby actually reducing the autonomy of

the decision maker. Collste *et al* (1999) also state that, whilst a DSS can assist in boosting the autonomy of a decision maker, that is by no means guaranteed.

As Silver (1988; 1991) points out, the finite processing limitations of any decision support system lead to restrictions on the decision-making abilities of the user. Whilst these may be of little or no consequence in many cases, few developers stop to consider that they are directly impacting the cognitive structures of the users of their systems. Indeed, the paternalistic approach to decision support is alive and well, as evidenced by the following from a recent text on data warehousing (Craig, Vivona, & Bercovich, 1999, p.321):

> *Standard reports can be an asset to an organisation because they limit the choice for users when it comes to researching decisions. By telling the users what they should be looking at, the designer of the standard reports removes the burden of deciding what is important and what is not.*

Clearly, this attitude is one of paternalism and if adopted, abrogates the autonomy of the decision maker to determine for themselves "what is important and what is not." In such a situation, it would be feasible to argue that the developers shoulder the ethical responsibility for the consequences of decisions based upon the output of their systems. That is, the locus of responsibility for decisions made by the decision maker shifts to the developer.

This is not to say, however, that all systems that adopt a paternalistic approach to decision support are necessarily unethical. It may well be that the user wants this level of support and structure. Having someone else decide for you what is important and what isn't removes a lot of the complexity from a decision situation. However, the relinquishment of the right to autonomy must be the prerogative of the user, never a result of a unilateral decision of the developer. A paternalistic system developed to meet the user's needs respects and maintains the user's autonomy only if the user's decision to adopt such an approach itself meets the criteria for autonomy.

Justice

Beauchamp & Childress discuss justice from within a health context, and look at issues of equality of access, fairness, and allocation of health resources. Of the four principles, justice is perhaps the least relevant to decision support, particularly individual as opposed to organisational decision support. That being said, the themes of equity and fairness do have implications for decision makers, particularly when these decisions have a social or strategic policy making aspect to them. By extension, especially for active and normative decision support, issues of social justice, equity and fairness are relevant. For example, the nature of decision support provided will have an impact on the role that stakeholders other than the decision maker, if any, play. The broader social issues of technology use referred to in the section above on Ethics and Information Technology are also encompassed by the concept of justice.

Concluding Comments

Ethics is not a side issue for DSS development. It should pervade every aspect of development, deployment and use of the system, covering not only professional conduct on the part of the developer, but consideration of the impact the system has on the user, and other stakeholders in the system and the decisions made relying upon it.

The paucity of research in this area, however, is disappointing and problematic. Apart from a small number of conference papers, there is almost no academic consideration of the issue. It is not enough to rest on the laurels of the work done on ethics in IT in general – the nature of DSS development and use mean that there are significant differences to other IT systems. First, the development process itself is quite different, being much more collaborative in nature. The relationship between the user and developer is generally much closer than in other systems development. The scale and length of projects is also generally much smaller and more 'intimate'. This is so that the second major difference can be catered for, that is, that the nature of DSS and the development process is a much more invasive intervention than for other systems. The collaborative nature of the development process and the intrusiveness of the intervention mean that there are a number of significant similarities between DSS developers and clinicians.

Neither can we rest upon the laurels of the ethical work in medical DSS, a result of the more advanced status of debate on ethics in medicine. The discussion has not translated across to DSS at large, as that work has been carried out by medical researchers and published in medical journals such as *Methods of Information in Medicine*. At the very least, if there has been any interaction between DSS and medical researchers, it has not translated into published research on ethics within the business DSS domain.

Beauchamp & Childress' principles of bio-ethics are not the only possibly useful framework for tackling DSS ethics. However, given the strong similarities between clinicians and DSS developers, we believe that the four principles of beneficence, non-maleficence, autonomy and justice provide insight into the many ethical aspects of DSS. This paper, therefore, is something of a call to arms for researchers to develop the area with both philosophical, as well as empirical treatment.

It is important that further work be done. Whilst we have specifically addressed *personal* decision support systems in this paper, we can't see any significant reason why the ideas here could not be broadened to include the other kinds of decision support mentioned in the introductory section – at the very least, it bears further investigation.

In terms of future research, we propose a four phase agenda, with each phase consisting of one or more research projects in itself. First is a need to refine and develop the framework conceptually. A more rigorous literature review, including a stronger input from the field of ethics is needed. Further conceptual development will allow expansion of the framework to include the other types of DSS described in the introduction. This will lead to a refined conceptual model of an ethics of decision support.

The second phase is to canvas input from DSS developers. Professional input will ground the model in the kinds of issues that are faced in DSS projects, and provide extra face validity. The third phase will be to take this model and test it in situ. This third phase could test a number of hypotheses. We believe that DSS developed with these ethical principles in mind will be more successful from the perspective of user satisfaction and system use, than would otherwise be the case.

Finally, with an empirically refined and validated set of principles, there is a need to see these principles put into use in practice. Broadly speaking, there are two approaches. The first is to evangelise to existing developers. This can be achieved through a number of methods, such as presentations, lectures, seminars, articles in practitioner journals, training courses, and so on. The second is to educate up-and-coming developers to think about ethical issues. This

involves modification of course syllabi in universities at both the undergraduate and postgraduate levels.

Phase	Methodologies	Outcome
1. Conceptual Development	• Literature Review • Argumentation	Refined conceptual model
2. Empirical Validation / Refinement	• Focus Groups • Surveys	Professionally informed conceptual model
3. Application to Practice	• Case Study • Ethnography • Surveys	Ethical framework for professional DSS practice
4. Evangelism / Dissemination	• Training • Syllabus development • Seminars • Lectures • Trade press articles • Books • Etc.	Ethically trained and aware DSS developers

Table 1. Research Agenda.

This paper is therefore a call to action. Whilst we will be continuing work on this research agenda, this paper is also a call for a healthier debate amongst the practitioner and academic DSS community. As academics, we have a responsibility to initiate and foster this very important discussion.

References

Abbott, PA (2001), 'Ethical Considerations for Decision Support Systems: Panel Presentation on Quality, safety and ethical issues in the use of computers to advise on patient care.' MEDINFO2001, 10th World Congress on Health and Medical Informatics, viewed 4th June 2003, <http://www.openclinical.org/docs/ext/safetypanel/abbott.pdf>.

Arnott, DR, O'Donnell, PA, & Grice, M (1993), 'Judgement Bias and Decision Support Systems Research', in *Proceedings of the 4th Australasian Conference on Information Systems*, Brisbane, Australia, pp. 65-80.

Banerjee, D, Cronan, TP, & Jones, TW (1998), 'Modeling IT Ethics: A Study in Situational Ethics', *MIS Quarterly*, vol. 27, no. 2, pp. 31-60.

Beauchamp, TL, & Childress, JF (1989), *Principles of Biomedical Ethics*, 3rd edn, Oxford University Press, New York.

Berleur, J, Duquenoy, P, & Whitehouse, D (eds) (1999) *Ethics and the Governance of the Internet*, International Federation for Information Processing, Laxenburg, Austria.

Chae, B, Courtney, JF, & Paradice, D (2002), 'Incorporating an Ethical Perspective into Decision Support Systems Design', in *Proceedings of the IFIP TC8/WG 8.3 Open Conference on Decision Making and Decision Support in the Internet Age (DSIAge2002)*, Cork, Ireland, pp. 136-153.

Collste, G, Shahsavar, N, & Gill, H (1999), 'A Decision Support System for Diabetes Care: Ethical Aspects', *Methods of Information in Medicine*, vol. 38, no. 4-5, pp. 313-316.

Conger, S, & Loch, KD (1995), 'Ethics and Computer Use', *Communications of the ACM*, vol. 38, no. 12, pp. 30-32.

Courbon, J-C, Grajew, J, & Tolovi, J (1978), *Design and Implementation of Interactive Decision Support Systems: An Evolutive Approach*, Technical Report, Institute d'Administration des Enterprises, Grenoble, France.

Craig, RS, Vivona, J, & Bercovich, D (1999), *Microsoft Data Warehousing: Building Distributed Decision Support Systems*, Wiley, Toronto, Canada.

Dowling, C (2001), 'Intelligent Agents: Some Ethical Issues and Dilemmas', in *Proceedings of the 2nd Australian Institute of Computer Ethics Conference*, Canberra, Australia, pp. 28-32.

Eom, SB, & Lee, SM (1990), 'A Survey of Decision Support System Applications (1971-1988)', *Interfaces*, vol. 20, no. 3, pp. 65-79.

Fox, J (1990), 'Automating Assistance For Safety Critical Decisions', *Philosophical Transactions of the Royal Society of London - Series B: Biological Sciences*, vol. 327, pp. 555-567.

Fox, J (1993), 'Decision Support Systems as Safety-Critical Components: Towards a Safety Culture for Medical Informatics', *Methods of Information in Medicine*, vol. 32, no. 5, pp. 345-348.

Goodman, KW (ed) (1998) *Ethics, Computing and Medicine: Informatics and the Transfornation of Health Care*, Cambridge University Press, Cambridge, UK.

Gray, P, & Watson, HJ (1998), *Decision Support in the Data Warehouse*, Prentice Hall, Upper Saddle River, New Jersey, USA.

Jelassi, MT, Williams, K, & Fidler, CS (1987), 'The Emerging Role of DSS: From Passive to Active', *Decision Support Systems*, vol. 3, no. 4, pp. 299-307.

Johnson, DG, & Mulvey, JM (1995), 'Accountability and Computer Decision Systems. (Ethics and Computer Use)', *Communications of the ACM*, vol. 38, no. 12, pp. 58-64.

Keen, PGW (1980), 'Adaptive Design for Decision Support Systems', *Data Base*, vol. 12, pp. 15-25.

Keen, PGW (1987), 'Decision Support Systems: The Next Decade', *Decision Support Systems*, vol. 3, no. 3, pp. 253-265.

Kimball, R, & Merz, R (2000), *The Data Webhouse Toolkit*, Wiley, New York.

Laudon, KC (1995), 'Ethical Concepts and Information Technology', *Communications of the ACM*, vol. 38, no. 12, pp. 33-39.

Lucas, R (2001), 'Why Bother? Ethical Computers - That's Why!' in *Proceedings of the 2nd Australian Institute of Computer Ethics Conference*, Canberra, Australia, pp. 33-38.

Mallach, EG (2000), *Decision Support and Data Warehouse Systems*, Irwin McGraw-Hill, Boston, USA.

Mason, RO (1995), 'Applying Ethics to Information Technology Issues', *Communications of the ACM*, vol. 38, no. 12, pp. 55-57.

Meador, CL, & Ness, DN (1974), 'Decision Support Systems: An Application to Corporate Planning', *Sloan Management Review*, vol. 15, no. 2, pp. 51-68.

Miller, RA, & Goodman, KW (1998), 'Ethical Challenges in the Use of Decision-Support Software in Clinical Practice'. In *Ethics, Computing and Medicine: Informatics and*

the Transformation of Health Care, ed. KW Goodman, Cambridge University Press, Cambridge, UK, pp. 102-115.

Power, DJ, & Kaparthi, S (1998), 'The Changing Technological Context of Decision Support Systems', in *Proceedings of the IFIP TC8/WG8.3 International Conference on Context Sensitive Decision Support Systems*, Bled, Slovenia, pp. 41-54.

Raghavan, SA, & Chand, DR (1988), 'A Perspective on Decision Support Systems', *Computer Science and Informatics*, vol. 18, no. 1, pp. 7-36.

Schein, EH (1987), *The Clinical Perspective in Fieldwork*, vol. 5, Sage University Papers Series on Qualitative Research Methods, Sage Publications, Thousand Oaks, USA.

Shim, JP, Warkentin, M, Courtney, JF, Power, DJ, Sharda, R, & Carlsson, C (2002), 'Past, Present, and Future of Decision Support Technology', *Decision Support Systems*, vol. 33, no. 2, pp. 111-126.

Silver, MS (1988), 'On The Restrictiveness of Decision Support Systems', in *Proceedings of the Organisational Decision Support Systems: Proceedings of the IFIP WG 8.3 Working Conference on Organisational Decision Support Systems*, Lake Como, Italy, pp. 259-270.

Silver, MS (1991), 'Decisional Guidance for Computer-Based Decision Support', *MIS Quarterly*, vol. 15, no. 1, pp. 105-122.

Snapper, JW (1998), 'Responsibility for Computer-Based Decisions in Health Care'. In *Ethics, Computing and Medicine: Informatics and the Transformation of Health Care*, ed. KW Goodman, Cambridge University Press, Cambridge, UK, pp. 43-56.

Suvachittanont, W, Arnott, DR, & O'Donnell, PA (1994), 'Adaptive Design in Executive Information Systems Development: A Manufacturing Case Study.' *Journal of Decision Systems*, vol. 3, pp. 277-299.

Thomsen, E (1997), *OLAP Solutions: Building Multidimensional Information Systems*, John Wiley, New York.

Walsham, G (1996), 'Ethical Theory, Codes of Ethics and IS Practice', *Information Systems Journal*, vol. 6, no. 1, pp. 69-81.