Economic, Legal, and Social Implications of Information Technology

In the March 1988 issue, the Editor regretted his "inability to stimulate a flow of articles and manuscripts relating to the broader societal issues of the Information Age." This note is a response to his comments. It considers why IT researchers and practitioners have not contributed sufficiently to the social implications debates, and how this deficiency might be rectified.

Conventional Approaches to "Social Implications"

Many years ago, MIT's Joe Weizenbaum, frightened by the simple-mindedness with which his Eliza experiments had been misunderstood, expressed serious reservations about artificial intelligence's implicit assumptions and directions. He came to be seen as an adversary of those who had been colleagues, as a whistle-blower, and that particularly dangerous opponent, a convert. Whether he intended to, he went out on an academic, intellectual, even spiritual limb. His work has had significant effect on people outside the computing-related disciplines, but is for the most part rejected or ignored by those inside them. **The adversarial approach** to the social implications of IT is not to be recommended — because of the intellectual isolationism it can lead to and its inherently limited effectiveness.

Rather than becoming converts to the "computing is dangerous" school of thought, some IT researchers and professionals have developed a double life. They have continued to practice their main disciplines, but have undertaken parallel activities on unrelated social implications matters. This is akin to the international arms-dealer who commits some of his spare time to the local Boy Scouts troop (an activity which I am not in any way intending to belittle). This **"independent roles" approach** to the social implications of IT has been apparent among scientists and professionals of all kinds in relation to nuclear arms, and more recently SDI.

It is inherent in the adversary model that the individual abandons his discipline, and/or his discipline abandons him. The second model, on the other hand, enables the person to retain his credibility in his own area of expertise. However, this is achieved by ensuring that one's professional and "social conscience" realms are not too tightly intertwined, and it requires considerable feats of ingenuity, persuasion, even sophistry, to inculcate an aura of authority in areas where one has limited expertise.

A development on this theme is **the dual-specialist approach.** Stanford's Terry Winograd and NYU's Ken Laudon come to mind, both of whom have undertaken fairly formal studies of particular social implications matters, while maintaining their other, more mainstream activities in AI and MIS respectively. As long as you maintain respectability in the mainstream, your proclivities for soft, socially responsible interests will not do you too much harm, and you can speak with authority in both of them.

Another approach is to confine your comments on social implications to areas in which your authority is well-recognized. Well-known computer scientists David Parnas of Queen's College, Kingston, Ontario and Brian Cantwell Smith of Stanford have not only undertaken research and written and spoken on SDI, but have also refused to use their particular knowledge and skills in its support. This approach has the benefits of considerable impact. It loses the individual some of his research grants, but protects his personal integrity and public credibility. But relatively few people have sufficient standing to use this **authoritative renegade approach,** and for the vast majority of us it would just be an excuse for inaction.

It is a human characteristic to seek out the company of others of like mind. As a result, special-purpose organizations and pressure groups exist, including the major players listed at the end of this note. Valuable though the efforts of these organizations are, their existence underlines the manner in which thinking about implications is being divorced from research in and practice of the technology.

Of course, an approach is available whereby we may excuse ourselves entirely. All we need to do is define the observation and criticism of the

IT disciplines as someone else's problem. After all, we have Maggie Boden of the University of Sussex (a psychologist/philosopher) to comment on the artificiality of artificial intelligence, and MIT's Sherry Turkle (the sociologist who wrote *The Second Self*) and Alan Bolter (the philosopher who wrote *Türing's Man*) on humans' self-conception in the computer age. Sociologists Jim Rule of SUNY at Stony Brook and Gary Marx of MIT have studied IT as a weapon of social control, and Rob Kling of California at Irvine has researched and written on the whole gamut of social implications.

The avoidance approach claims that self-criticism by the IT disciplines is either unnecessary or too hard, and should be left to outsiders. Apart from being a moral "cop-out," the weakness of this approach is that outsiders can never fully appreciate the developments in specialized fields unless they have the assistance of insiders. External criticism would therefore be poorly aimed and ineffective unless at least some IT specialists collaborated with critics of their discipline.

"There Must Be Another Way" Department

I don't believe that the adversary, independent-roles, dual-specialist, authoritative renegade and avoidance approaches represent a sufficient set of models. I believe that **all researchers and professionals must regard the implications of their work as part and parcel of their research in and application of IT. Consideration of implications needs to be integrated, not segregated.**

A fundamental barrier is that we continue to subscribe to a vague image of scientist as saint, or as virgin. We pretend that Copernicus didn't hide his discovery, that Kepler wasn't a mystic, that Galileo wasn't a trickster and an intellectual thief, and that Benjamin Franklin wasn't (among many other things) a politician. We represent our work as being value-free. Or, at the very least, we go to excruciating lengths to convince ourselves that we have made best efforts to rid our work of the very values that might make it relevant.

Naturally I am not arguing against care, precision, external standards, and controls, above all in the design of research and in the bodies of papers. What I am attacking is the pretense of pristine purpose. The introductions to our papers should recognize the full-blooded relevance of the topic, and our implications sections and our conclusions should shed the veil of scientific disinterestedness and admit to motivations and to concerns.

Another barrier preventing the integration of implications with applications thinking is that the social implications area has the smells of altruism and intellectual softness about it. "Real programmers don't worry about social implications" is an easy and realistic catch-cry. It's a little like the Moral Rearmament movement — it's a good thing of course, like motherhood and apple pie, but you have to be a bit of a fanatic, and maybe even a bit of a wimp, to pursue it.

About three years ago, I assumed the chairmanship of the Australian Computer Society's Social Implications Committee. I was concerned at the narrowness of the Committee's Terms of Reference, because it seemed that it must inevitably comprise "do-gooders," putting their fingers in dikes, and acting as the moral conscience of the IT professions. To address the wimpy image, I proposed that the various dimensions of IT impact should be considered together, in an integrated manner.

In 1986, the Society agreed to enlarge the Committee's Terms of Reference and it became the Economic, Legal and Social Implications Committee (ELSIC). During the next two years, the so-called "Australia Card" scheme consumed a considerable proportion of the Committee's time. The Australian government proposed to introduce a national identification scheme, as a basis for a local version of the often considered but always rejected American idea of a "national data center." Civil libertarians from all segments of society attacked the scheme because of its, in part describable and in part imagined, social implications.[1]

ELSIC expressed concern about the social implications, but only to the extent that was necessary to argue for careful consideration of the proposal. The focus of submissions to government and the various review committees was on assessing the scheme's technical base and its eco-

[1] See Clark, R.A. "Just Another Piece of Plastic for Your Wallet: The Australian Card," Prometheus (5:1), June 1987. Republished in *Computers & Society* (18:1), January 1988 (the quarterly journal of ACM SIGCAS).

nomic merit. ELSIC was therefore able to skirt the more subjective aspects of the debate and speak with authority on matters susceptible to more precise analysis. **The conventional topic of "social implications of computing" is too narrow. Economic, legal, and social implications of information technology must all be considered together, to enable the various factors to be seen in perspective.**

The Moral Issue

To what degree does a technologist have moral responsibility for the uses to which his technology is put? Professional codes of ethics frequently require a member to place the public's interests above his own and to avoid nasty uses of technology. But the meaning of the inevitable judgmental words (like "nasty") are dependent on context. Agreement on the applicable standards is difficult in time of peace and seriously problematical during periods when threats of terrorism or invasion exist. Such threats may be real or imagined — but at the time, with propaganda emanating from the "other side," and managed information flows on "our side," it is by definition impossible to tell the difference between reality and imagination. And anyway, codes of ethics are seldom enforced, or even remembered.

Researchers and practitioners in IT are dealing with a powerful tool. We have clearly before us, as archetypal anti-heroes, the nuclear physicists of the 1920s and '30s, who hid their heads in the sand and comforted one another that the implications of their work were unforeseeable. Or that if they were foreseeable, then their work was a-moral, not immoral, because it was other people who would decide on the uses to which nuclear energy would be put. Should we indulge ourselves in our own modern versions of this avoidance approach, or risk our scientific detachment by becoming involved? To co-opt Kennedy's inversion of the question, if we believe our own hype that IT will have enormous impact, can we afford *not* to risk our scientific detachment and become involved?

I believe that **the moral responsibility of any professional must at least extend to an honest attempt to ensure that public debate is informed. Further, since the subject-matter is often obscure, the professional's role extends to ensuring that debate takes place.** At times, that step will drag the person into the role of a protagonist or antagonist. Clearly in those circumstances we are no longer entitled to claim the mantle of technologist-saint; but then aren't we all parents, voters, patriots, and even civil libertarians and ideologues as well?

If MIS academics and thinking practitioners are to stimulate and support public debate, we must make some changes in our conventions of behaviour. We do not need to compromise the precise, careful, scientific manner in which we undertake and report on our research. But we must stop sanitizing our introductory remarks, and instead draw attention to the real importance of the topic we are dealing with. And the closing sections of our papers must not be confined to "implications for further research," but must also directly address "implications for people."

It is important that researchers do not restrict their writing to the serious academic journals on which our career prospects and peer approval depend. We must also report the nature and implications of our work in more populist publications, in order to reach the great majority of practitioners whose professional reading does not reach the heights of the *MIS Quarterly*.

Like many others, I missed the Editor's "call to arms" in the March 1986 issue. As a result, although I have regarded this as one of the most important journals in which to read and publish technical and managerial articles, I have never even considered looking in it for papers on technological implications. I therefore very much welcome the Editor's intentions.

However, I suggest that, by broadening the topic from mere "social implications" to "economic, legal and social implications," we may be able to overcome our reticence about thinking, talking and writing on the consequences of our work and to integrate the consideration of implications with applications.

Roger Clarke
Department of Commerce
Australian National University, Canberra

Organizations and Current Chairmen

IFIP TC9
IFIP Technical Committee TC9
(Relationship Between Computers and Society)
of the International Federation for Information Processing
Prof. Hal Sackman
13609 Bayliss Rd.
Los Angeles, CA 90049

IFIP TC9 Working Group WG9.1
(Computers and Work)
Prof. Klaus Fuchs-Kittowski
Humboldt Universität Berlin
Sektion Wissenschaftstheorie und-organisation
Unter der Linden 6, psf 1287 DDR-1086
Berlin G.D.R.

IFIP TC9 Working Group WG9.2
(Social Accountability)
Dick Sizer
26 Avenue Rd.
Farmborough Hants GU14 7BL U.K.

ACM SIGCAS
Special Interest Group on Computers and Society of the (charmingly named) Association for Computing Machinery
Ron Anderson
13221 Lake Point Blvd.
Belleville, MI 48111

CPSR
Computer Professionals for Social Responsibility
P.O. Box 717
Palo Alto, CA 94302

BCS SIC
Social Implications Committee
of the British Computer Society
13 Mansfield St.
London W1M 0BP U.K.

ACS ELSIC
Economic, Legal and Social Implications Committee of the Australian Computer Society Inc.
Roger Clarke
Department of Commerce
Australian National University
Canberra ACT 2611