The Ethics of Research into Invasive Technologies

Bob Anderson

Technical Report EPC-1991-107

Rank Xerox Research Centre
Cambridge Laboratory
61 Regent Street
Cambridge CB2 1AB

Tel:+44 1223 341500
Fax:+44 1223 341510

The Ethics of Research into Invasive Technologies

Bob Anderson

RXRC, Cambridge Laboratory

Author's address:

Rank Xerox Research Centre,Cambridge Laboratory,
59-61,Regent Street,
Cambridge, CB2 1AB.

Email: Anderson.EuroPARC @ Rx.Xerox.Com
Tel: +44 223 341507
Fax: +44 223 341525

The last temptation is the greatest treason:
To do the right deed for the wrong reason.
(T.S. Eliot. *Murder in the Cathedral*)

--------||--------

The last temptation is the greatest treason:
To do the right deed for the wrong reason.
(T.S. Eliot. *Murder in the Cathedral*)

--------||--------

Canon 5

EC5.2 An ACM member, whenever dealing with data concerning individuals, shall always consider the principle of the individual's privacy and seek the following:

To minimise the data collected.
To limit authorized access to the data.
To provide proper security for the data.
To determine the required retention period of the data.
To ensure the proper disposal of the data.
(*ACM Code of Professional Ethics*)

--------||--------

Preliminaries

Friends tell me my writing is sometimes a bit opaque. In particular, it is often difficult to see which point I am pursuing. This leaves them puzzled why I belabour the trivial and the obvious as I do, and why I am so fascinated with obscure examples and apparently irrelevant diversions. Given the possibility of such obscurity, I want to spend a few moments laying out my overall objectives and the general approach to the topics covered in this paper. In the spirit of reader-centred writing, my hope is to prevent needless confusion and irritation.

To begin with: the important matters of style. This is not a research report in the normal sense nor is it a summary review of the literature in the fields it covers. I am primarily interested in arguments not findings and, in particular, with the forms which arguments take, not with classifying writers or writings according to the arguments which they use. Most importantly, this paper is a way of clarifying *my thinking* about these matters and an invitation to others similarly engaged to join me in dialogue. In no sense is it

meant to be an introduction to research on 'invasive technologies' nor a quick summary of the state of play for the interested observer.[1]

Second, since I am interested in arguments, what especially engages me here is the route that can be taken through a set of propositions. How can one argue one's way from position to position in a consistent and coherent fashion? The questions I raise then are about tracking an argument from beginning to end and what it might take to have a trackable, coherent argument. What I am very specifically not interested in are design decisions and their implications.

For me, this outlook imposes a number of features upon my discussion.

The issues are conceptual not substantive:

The discussion focuses on how we are to assess the significance (in this case *ethical* significance) of any findings we might have. It does not seek to define design questions nor make judgements between research strategies and policies in the light of whatever information is to hand.

The character of invasiveness is irrelevant to this discussion

The determination which, if any, technologies should be designated "invasive" or as potentially so, is left open. Simply that they are seen as such warrants my discussion of them since it is arguments about the ethics of research on them that concerns me. Suffice it to say that a number of research lines currently being carried out at Xerox PARC, and Rank Xerox EuroPARC, Bellcore, The University of Toronto and elsewhere have been dubbed with this epithet.[2] Of particular concern appear to be research activities which focus on the exploration of multi-media shared spaces and of mobile location devices (e.g. the Active Badges originally developed by Olivetti Research Labs in Cambridge (Want 1992)). As far as I am concerned, it is an interesting *sociological* issue why these technologies should be the cause of such concern (and not, for in-

---

1 There is a serious literature which is summarised in Dunlop and Kling (1991) and Mackay (1991). A more journalistic account can be found in Lacayo (1991).

2 For PARC, see Harrison et al.(1990) and Weiser (1991). For EuroPARC, see Gaver et al. (1991) and Newman et al (1991). For Bellcore, see Fish (1990). For Toronoto, see Mantei et al. (1991)

stance, CSCW tools, teleworking, or Computer Aided Management).

The general issue of social consequences and their ethical implications has, of course, been widely discussed. Such discussions generally take one of two forms. Much has been said about the issue of 'dataveillance' (Clarke 1989) and the ways in which computational technologies can facilitate organisational and governmental monitoring of individual activities. With regard to the technologies I have in mind, not only are parallels drawn between the potential dangers arising from systems of electronic fund transfer, secondary data base accessibility and the like, but the technologies themselves are held to be, if anything, more pernicious since they seem somehow to be more direct in their impact. They seem to have implications for our conception of ourselves as social individuals as well as our identities as citizens, employees, members of work teams and so on. Second, many authors (Ladd 1989, Ellul 1989, Pullinger 1989) have commented that the very character of these technologies is rendering our concepts of rights, obligations, and responsibilities otiose. The social implication which is drawn is that we need to develop either a modified or a completely new set of cultural norms which is specifically able to take account of these technological innovations. The recognition of this is, of course, what underlies the ACM Code of Practice which I cited at the beginning.

Since I am not concerned with deciding which are and which are not 'invasive technologies', one set of issues I do not address directly is that surrounding the vexed notion of privacy. Although the potential impacts on the individual and on society of restrictions on privacy might feature in arguments about these technologies, the concept is rarely, if ever, explicated. Ways of drawing the line between the public and the private, therefore, remain familiarly intuitive. Since all parties to the discussion operate in the same way — the differences between them seeming to reside in the character of their intuitions — there is no obvious reason for my discussion to take that particular turn-off.

There is, also, a further and much bigger issue buried here the dimensions of which I can only guess at. This is the changing moral character of modern large scale economic and other kinds of organisations, and hence the range of ethical principles to which we hold them accountable. Certainly, it would be as naive to suppose that there are no relevant commercial considerations to their increasing concern with customer satisfac-

tion, quality of working life and environment, good citizenship and the like, as it would to believe that commitment to these objectives was always being driven along by an obsession with ‘the bottom line’. Quite how research plays into and is effected by these trends is hard to determine. All one can say is that strictly ethical considerations as well as economic ones are increasingly likely to be under consideration.

In the case we are discussing, along with this set of confusions are introduced other cross-currents. North American and European high-technology Companies are under threat from the Far East. Part of the response has been to place emphasis on a core competence — successful Research and Development. Researchers have, therefore, become drawn into the mainstream of company activities, but are not necessarily more influential. More is asked of research: and that more is couched in terms of products, profitability, time to market etc. With this desegregation, it may become more and more difficult to maintain potentially antithetical value systems (such as those of the research community and the market place) within one organisation. The motivations for research choices — what determines an important or interesting problem, say — may well be moving. At the same time, as we have already mentioned, clear attempts are being made to treat firms as socially responsible agents and not as the asocial profit maximisers of Economic theory.

These are deep issues on which we as a community of researchers should reflect. In this brief discussion, I do not expect to get beyond sketching a couple of initial clusters of concerns and perhaps encouraging the opening up of the debate up a little.

What comparisons are relevant?

Since I am discussing what we want to say about a particular cluster of ways of working, i.e. research, a great deal can be gained by comparison to similar practices where ethical arguments are agreed to be relevant. Indeed, the force of this view is apparent from the fact the research community is itself constantly seeking to make appropriate comparisons (Mackay 1991).

The upshot of the foregoing stylistic features is that it is extremely unlikely anyone will be able to extract a definitive conclusion from my reflections. Indeed, I should feel deeply unhappy were this not the case. The sheer difficulty of arriving at definitive conclusions in this area is probably the only real conclusion that can be drawn. This

does not dismay me though because, to repeat, I am interested in the form of the arguments, the grounds on which one might argue, not in arguing my way to any particular conclusions. To invoke a useful comparison: this is more a discussion of the conventions for map-making and wayfinding than it is a choice over the quickest or easiest route from A to B.

I hope I have said enough about what I am not going to do. Let me now try to crystallise what I do want to achieve. Basically, I want to argue that invoking a conventional ethical framework, i.e. standard utilitarian principles, as the basis for determining the ethical nature of research into 'invasive technologies' is not as straightforward as might appear.[3] To do it successfully, and that means to have an adequately defensible utilitarian argument for or against this kind of research, requires much more clarity about the form and the substance of the arguments in play. My hope for the present paper is to contribute to that clarification.

---

3 Utilitarianism is not the only set of principles which might be invoked. One could well imagine intuitionist and transcendental ones which would do as well (or badly). To Hemmingway is attributed one well known intuitionist principle: moral is what you feel good after. Kant's transcendental principle is equally widely cited: 'Act only on that maxim through which you can at the same time will that it should become a general law.' The point is that utilitarianism in its standard consequentialist form typically provides the framework within which this and related arguments proceed.

The Form of the General Issue

The juggernaut of technology

Let us begin with the simplest and perhaps most common case. This is a proposition about the inherent logic of technological innovation. The technology, we are told, is remorselessy moving in a particular direction. If ‘we’ do not do the research necessary to control, humanise or in some other way shape its development, others will. And they may be a great deal less scrupulous about or effective at ‘getting the technology right’. The contrasting argument has a very similar form. It goes something like this. These technologies are irredeemably pernicious. Nothing we can do will remedy this and any work we do carry out is quite likely to make matters worse. We should, therefore, anathematise all work on them. I will, for the most part, consider the first version, although the considerations I raise apply to the second as well.[4]

I find these claims intriguing for a number of reasons, the least interesting of which is that I do not know whether they are true. That is, I do not know what hard evidence one could muster to show that any single group was more likely than any other to mitigate the supposed effects of these potentially invasive technologies. However, as I say, that is the least interesting aspect of the question. Much more enticing is the *kind* of argument which we are being offered as a rational justification for this type of research and the counter-considerations which are relevant to it.[5] At one level it looks to

---

4 There are three issues which although very important in the broader context, are not directly germane to my concerns here. The first is the supposition that the potential invasiveness of these technologies makes them undesirable *sui generis*. In my view, this is an open question. However, for the purposes of this argument, I will assume these technologies are, in some sense, suspect. The second concerns who ‘we’ are supposed to be. In the contexts in which I have encountered the argument, this might be rendered variously as ‘we’ the HCI community, ‘We’ in Computer Science, ‘we’ in Xerox.....Third there is the horatory aspect. Presumably, we would want to distinguish between demands that no-one should engage with these technologies, and the justifications which might be offered by any individual for selecting some other interesting set of problems to work on. It is arguments of the former type that interest me.

5 The rational character is crucial, as we will see. An entirely different rhetoric, emotivism perhaps, would work in very different ways and have to be assessed against other criteria. The search for a rational basis for ethical choice derives from the desire to premise ethical principles on right thinking (to use an archaic term) rather than whim of fancy. The latter are envisaged as a Pyhrronean ethical relativism where not only do we have to tolerate differences of grounding between individuals but additionally no individual need feel obligated to be consistent in his or her decision making.

be a peculiar kind of special pleading. At another, it has all the hall-marks of a standard consequentialist ethical principle.[6] The question I want to take up now is, simply, which is it? Or, to put things more clearly: can we justify research on invasive technologies along lines similar to the consequentialist justifications which might be offered for research into genetic engineering or chemical warfare, for instance? And, if the parallels with these cases fail, is special pleading all that we have left?

I raise these questions for two reasons. First, I really do not know what to say here. I do not have a worked-out response and am pulled both ways.[7] As I say, my hope is that through working through (some of) the relevant considerations, I might be better able to clarify my own thinking and hence be able to make a judgement not just about this case, important though it is, but in other similar cases. Second, and perhaps as important as the first, I am not sure that these issues are currently being given much attention in the discussions of, for example, CSCW tools and technologies, multi-media research, groupware, interconnectivity and the like.[8] Emphasis is, naturally, being given to the positive, enabling potential of systems and to identifying and minimising potentially negative side effects, as if enumerating both of these would make deciding on the ethical correctness of the research unproblematic. I am not sure

---

6 Consequentialism seeks to provide a rational justification for a line of action by weighing relative outcomes. This is usually done by trying to evaluate the contribution which the activity in question might make to the commonweal or stock of public happiness (utility). Actions which add to the stock are justifiable: those which reduce it are not. Since a method for weighing outcomes is needed, some form of cost/benefit analysis is usually envisaged.

Consequentialism is not the only rational ground on which one might seek to defend the ethics of research either. One might appeal, for instance, to rights of the individual to pursue knowledge for its own sake, the obligations one has to one's employer, the interests of the state, and so on. Since the current case is not explicitly couched in these forms, I will not consider them here. Alternatively, one could quarrel with the whole rationalist approach and demand a new set of principles altogether better suited to exposing the "unreason" of this technological age. This is roughly what Ellul (1989) does.

7 One definition of a genuine *philosophical* problem could be that when faced with it, we do not know what to say.

8 The ethical considerations I have in mind here are not the same as the supposed social implications such as 'de-skilling', the potential for loss of control of personal information, the shifts in the distribution of knowledge and power said to be consequent upon their introduction into settings. Such implications might well figure as components within a consequentialist argument. As I mentioned earlier, these and many related issues are widely debated (Dunlop and Kling 1991). What I am concerned with is how to go about evaluating the 'Trust me, I'm a good guy' kind of blandishment which often passes for argument.

that they do. Let me make it quite clear, though, that I am not saying this kind of research is illegitimate. What I have my eye on is the claim to legitimacy encapsulated above.

The Consequentialist Frame of Reference

Discussions of the ethics of particular lines of research usually take place in relation to such contentious matters as whether it is permissible to cause pain to some human or non-human subject in order to be able to relieve wider suffering; the rightness of working on life-extending therapies in an over-populated world; the relative or absolute value of life; the rights associated with persons; and, of course, the applicability of the concept of personhood to, for example, embryos.[9] Now, on the surface at least, there are direct parallels between these cases and that of invasive technologies, even if we think the issues engaged in with the former are weightier, more significant, or, somehow, more serious. In each case, there is usually a clear intention to decide matters in terms of the net increase in the stock of public benefit by determining *a priori* the consequences of contrastive lines of action. Thus, the argument might run for our example, in the long run we will all collectively be better off by having these technologies even though in the short run, some of us may not. They will reduce costs, increase the availability of needed services, and have other beneficial outcomes. But do these parallels go any deeper? Are the cases isomorphic at every point? Can the arguments can be mapped from one domain to the other with no serious emendation? And if the parallels fail, what then?

The effectiveness of the consequentialist approach turns on the methods applied to weigh outcomes. This creates several levels of issues to be considered (Hare 1981). First there is the level at which we encounter the moral dilemma and find ourselves unable to see what we ought to do — what Hare calls the ‘intuitive level’. If we were not unable to see what to do, we would not have an *ethical* problem. The next level is the “critical level” where we apply a method which involves, perhaps, reflecting on options to determine courses of action. Third, there is the ‘meta-ethical level’ at which we debate the justificatory basis of the principles on which the method is based. The

---

9 I have leaned very heavily on Julian Glover’s ‘It makes no difference whether or not I do it’ in Peter Singer’s (1986) collection in framing this discussion.

first two are what we might term the domain of *mundane ethics*. The third is the province of *philosophical ethics*. As I hope I have made clear by now, my concerns are entirely with the third level, namely the grounding of the method. The essential contestability of ethical matters derives from the possibility of this third level and the confusion between it and the other two.

Simple Consequentialism

The most straightforward approach to weighing outcomes is to consider the relative value and scale of effects. If we use the relative contribution which engaging in the action will make to the stock of human happiness, then our consequentialism is utilitarian in character. Utilitarian consequentialism is the simplest case, though not the easiest to operate. In addition, on one relatively uncomplicated reading, our initial claim has strong utilitarian overtones.

The aim of the method is to determine precisely what difference engaging in a certain line of action will make. If we hear our initial justification as saying that our involvement makes *all* the difference to how things will turn out, we have an empirical claim about our unique capacities and a further counterfactual claim about how things would turn out if we were not to be involved.

We now have to face the problem of calibrating metrics. We might, for instance, be perfectly willing to say that if 100 (or 1000, or 10,000) people have a rare, painful and lethal blood disorder from which they could be could be saved by the suffering of 1 person, then that suffering might be justified. We might be less inclined to say so where it is the *inconvenience* of the 100 people which is at issue. If we were to extend waiting time for operations on in-grown toe nails and thereby create resources for one new home-dialysis machine, would that be justified? Probably so. But where on the range on the range of effects from lethal to inconvenient, do our invasive technologies sit? Do we know how to develop common scales for such disparate cases?[10]

Similar considerations apply to what appears to be the mutual exclusivity of the

---

10 In very large measure, the answer to this depends upon what is being invaded by these ‘invasive’ technologies. The usual response is privacy. But as I keep saying, that is a very obscure notion.

options. Is it being claimed that if we are not involved in the research, there will be an otherwise avoidable disaster? Here, we might want to say the lessons from the favourite case of the mother and baby apply. If some catastrophe occurs during pregnancy or birth, can we justify saving one and thereby losing the other? Can we equilibrate the risk of trying to save both (and potentially losing both) against deliberately sacrificing one? And if so, which (ceteris paribus) should we opt for?[11] Are the (obvious) major negative effects on the researchers (and their associated junior and support staff) who are using these technologies in their prototype forms justifiable by appeal to long term gains?[12] We have to do it, no-one else can; and so we have to put up with the invasiveness.

Probably, though, we would not think of our case in these exclusivist terms. We are far more likely to say that our expertise is such that we will make a major but not vital difference. For the moment, let us disregard side effects such as the consequence of our involvement being seen to give a good (or in this instance bad) example. Now we have to make decisions about thresholds. If we make 90% difference to how things turn out, should we still do it? 40%? 10%?

The problem here has two aspects. The first is simply in measuring effects at any threshold. Can we definitively say what is gained or lost at each point as we slide from 90% difference to 10%? Is 'invasiveness' a partitionable concept like that? Then there is the matter of *our* involvement. Can we say what the effects directly attributable to us will be? If we can, how do they relate to the 'costs' which are incurred? Here the parallel seems something like the familiar 'free rider' problem. If there is a general exhortation, for example to reduce water consumption during an extended draught, how do we measure the effect of (and hence calibrate the justifiability of) any single person's action? Surely in a community the size of Southern California, for instance, the net loss to the collectivity of an individual's profligacy is so small as to be immeasurable? It makes no real overall difference. The net loss to the individual, however, might be enormous. From the individual's point of view, it must be rational

---

11 Alisdair MacIntyre has forcefully argued that this is a major misconstrual of what is at stake in this case. The moral choice is not between the mother or the child, but between acting to save one and not acting to save either (MacIntyre 1990).

12 For simplicity here I assume the case where the users are researchers.

to continue with old habits of water consumption.

Our case might fit an extended and inverted version of this. Research on these technologies is going to mean loss of privacy for the research team and associated staff, possible ridicule, and certainly some obloquy in the research community. On what grounds could we argue they ought to take it up? If their involvement makes some but not all the difference to the eventual shaping of the technology, at what point does our imprecation to them to join in the research cease to have ethical force? If we cannot define the sub-thresholds and indicate what happens as we pass them then we cannot evaluate the effect of an individual taking one course of action as against another. We cannot tell what difference *we* would make. And if we cannot tell what difference we would make, what reason would there be for us engaging in work that has major social costs even if we knew or believed that in the long run the collectivity would benefit?

These thresholds might apply in a different way though. For invasive technologies to be of maximal value to their users, they have to be ubiquitous (or nearly so). Everyone has to join in. We as researchers may have the skills to control the systems and the motivation to tolerate their negative features. So too might the non-technical staff with whom we work. They may well have the motivation to support the research as subjects, and as members of a working community be able to exercise pressure on the uses to which it can be put. But what of those who are not aligned with research in this way and whose involvement is 'volunteered' for them? As we step from one threshold to another, from research, to support in the lab., to the non-research environment, are any of these points especially significant? Are the effects felt at any of these points to be weighed differently? Further, can we expect these non-researcher subjects to have the same commitment to research outcomes we might researchers? If not, how and where are we drawing these lines? Here we bump up against the very broad issue of the ethics of social research.

All of the above considerations are well known defects of our simple (act) utilitarian construal of the problem. Would they be overcome if we were to adopt some other version such as *rule* utilitarianism? Under rule utilitarianism (Mackie 1977), we determine outcomes by looking not to the net sum of happiness or benefit deriving from *this or that* action but from considering how far 'the general performance of acts of each of these classes' would impact upon general happiness. However, the nub of the objec-

tions still remains. We seem to have no way to operate the calculus, no way to measure happiness. Or rather we have no way that is not itself subject to excessive dispute and hence unlikely to lead to consensus on the use of the method. This being the case, as an alternative strategy we could consider enumerating other indirect consequences which while we might not be able to say if they would lead to general happiness, nonetheless could be considered to be beneficial. These are side effects.

Consideration of side effects

I will run over just a few of the possible kinds of side effects which might be relevant to our consideration of research on invasive technologies.

Are there other beneficial research lines opening up?

While we might admit that these technologies make invasiveness possible, perhaps they also offer the chance to improve diagnostic systems, develop new and better educational aids, systems for the handicapped of various kinds, and perhaps ways of reducing waste, minimising costs and so on, in other areas. An argument for taking up invasive research might, then, be buttressed by appeals to the obvious gains to be made by applying that research in these and similar fields. In my experience, arguments of this kind are rarely openly deployed in the case of invasive technologies, though perhaps they could be.

Are there research opportunity costs?

If we take up research on invasive technologies, will we therefore forgo the chance to research on some other set of problems which is more socially useful or contribute more greatly to the net collective benefit? In other words, does doing *this* have the side effect of my not doing *that*? Taken on its own, this is not a decisive argument, since we could always have done something else. The issue is whether it is obvious what else we could have done and whether the benefit flowing from that research would be greater than from the work on invasive technology. Clearly this is a tricky set of issues to deal with because it leads us right back to the scaling problems we encountered just now. How do we align the benefits to be accrued from different research strategies?

What is our standing in the community?

If we take up these technologies, even though we have our own good reasons, will it serve to legitimise lines of activity which it might be preferable not to open-up? Even if benefit might flow from our actions, overall there might be a net negative effect because of the encouragement which they will be interpreted as providing for other, somehow less desirable, research. This side effect is a potentially ironic outcome of accepting the original claim. We do the research because if we don't someone else will ....., and find ourselves being used as the precursors for the very research we thought we were pre-empting. In that sense we might have to ask if there is a Gresham's Law of Research?[13] Perhaps, because we have entered the field, we will both create opportunities for work of this kind and help to shape the mores of others who have less sophisticated reasons for being in it. If researchers are attracted into a domain, does this necessarily lead to a devaluation of research and a degradation in quality, standards etc.?

Does the free rider argument apply?

One side effect could be that other researchers gain the benefits of technology innovations and lessons (e.g. about the social and organisational aspects of these types of technology) but at no cost or risk to themselves. They are 'free riders'. Is it fair that, as a researcher, I should benefit from the work of others who are engaged in these invasive technologies without any cost to myself? If I have the skills and opportunities, should I rely on others to pay the costs for me? The parallel case usually considered here is that of democracy. Should I benefit from the operation of democratic structures without being willing to take part in the processes which sustain them? The problem is that, as we saw earlier, free riding is, or appears to be, a rational strategy for the individual but irrational for society at large. Resolving that conundrum from within either act or rule utilitarianism still appears to be a long way off.

The point about arguments over side effects is that they are *side* effects and hence can

---

13 Gresham's Law states that bad money drives good from the market place and is a generalisation about the effect (originally of clipped coins and counterfeit notes) of unreliable or overly volatile investments in the money markets. Basically, people won't take risks in an excessively risky environment.

only help adjust the balance of judgement, not fix it. They will only be brought into effective play when we can reach no determinate conclusion on the main consequences. In as much as this is always the case for utilitarian versions of consequentialism, we should not be surprised to see side effects featuring strongly in discussions of the legitimacy of research on invasive technologies. What we should bear in mind though is that while they might be frequently used arguments, that does not make them strong ones.

Supererogation or Captain Oates' arguments[14]

Thus far we have been considering the evaluation of the moral force of the original claim in group terms. That is, we have been considering it from the point of view of a research community. There is another way of taking it though, one which is construed in individualist terms. While there is no obligation on 'us' as a group to embrace or reject this line of work, we as individuals might have an obligation to take the responsibility of sacrificing ourselves by doing this research. This obligation is a supererogatory one.

Virtuousness aside, what is the strength of this case?[15] Clearly it only works at all if the net sum of benefits to ourselves remains negative while the balance overall is positive, for if it were not there would be no talk of sacrifice. In a sense, then, it is a re-run of the first argument about the uniqueness of our contribution with our moral rectitude somehow equilibrating the situation and making it both rational and required for us to do the deed. But why should that be? Can we be required to be saintly? Once again we run into the difficulty that the arguments do not decompose downwards. It may be necessary, on utilitarian grounds, for there to be some saints. It may even be necessary for there to be some *research* saints. But it can never be necessary that I be one of them.

---

14 During the return leg of Captain Scott's fateful and unsuccesssful attempt in 1911 to be first to the South Pole, the team was reduced to eating their dogs and pulling the sledges themselves. Eventually, when they were almost out of food and still many days from their base, it was Captain Oates who made the decision to sacrifice himself and with the words "I am going out now, I may be gone a little while..." walked out of the tent and never returned. Sadly, eventually all the team perished on the ice.

15 Hare (1981) discusses supererogation only in terms of saintliness, that is those cases with the lowest 'acceptance-utility'.

On the other hand, we could turn the question around. Judith Jarvis Thomson (1986) asks about the case of a person who is (unwillingly) plugged into someone else with a rare kidney disease. The issue is not that the 'donor' is uniquely qualified to save this life and so should do so, simply that, for everyone, the moral obligation to save life where possible outweighs anything else. In such circumstances, although there would be no reason to demand any specific person should be so exploited, the moral obligation to be what she calls a 'minimally decent Samaritan' applies to everyone.[16] If we find ourselves faced with an occasion when we could behave as a minimally decent Samaritan, we should. The question is where one draws the line between the obligation to behave with minimal decency and the freedom to display goodness if one so chooses.

If we take these considerations to the arena of invasive technologies, then we have to ask whether the requirement on a set of researchers to take up a line of research, given that their involvement will make all the difference to the ways things turn out, can be justified by appeals to minimal decency. And, furthermore, how much of the social cost of carrying out the research (vilification etc.) they can be expected to bear.

In effect, the supererogatory interpretation of our original claim is an attempt to create a sense of professional responsibility for standards of research practice. As researchers, we are being asked to impose an ethic of professional responsibility on ourselves which will be enforced or enforceable by the rest of the community of researchers. The existence of this code would lead us to accept the obligation to research on these technologies should we have the opportunity and be competent to do so. But what of the reverse of this case — the equivalent of the foolhardy 'saint'? Would the same arguments apply if my *not* working on these technologies were to give them a better outcome. In other words, does this moral obligation of good professional practice regulate the omissions of others as well as their commissions? And what does that imply for other closely related areas?[17]

---

16 The "minimally decent Samaritan" is contrasted with the New Testament's Good Samaritan who felt obligated to risk his life, give up his goods, time, and money.....

17 Now we run up against the other side of utilitarianism, namely its libertarianism and the problem of demarcating how far individual freedom supervenes on collective responsibility.

Trojan Horse arguments

Trojan Horse arguments are interestingly different to both of the previous kinds but nonetheless somewhat allied to them. In essence, what is being argued is that research on invasive technology should be a means to some other (presumably more laudable) end. The technology (with all its properties) has now shifted to a side effect and hence been downgraded in status. So, while it might be accepted that this technology will sell and hence be justifiable in pure business terms, its inimical properties would make it illegitimate were it not for the 'leverage' (or whatever) it provides.[18]

Under what is probably the most plausible interpretation, this could be what the original justification was claiming. For this argument to hold, the consequentialist has to decide whether there is a net positive balance of effects if energy is channelled away from making the technology more saleable or cheaper or more broadly applicable and towards research lines which make it more difficult to use, more complex, more easily brought under the control of the user and so forth.[19] Here what is involved is the discovery of some means for balancing the obligation to ensure one's research is productive, cost-effective and contributes to return on assets (or whatever criteria are set) against the obligation to contribute to the scientific community and the promotion of the common good. And since that is where we started, we are still in need of some *arguments* to enable us to make choices in the situations we are faced with.

The difficulty derives not so much from the vagueness of these concepts nor from the impossibility of scaling outcomes, but from problems being set us by *this* technology. The cases usually discussed in this regard are those where whistle-blowing is the issue or commercial privacy. Should researchers reveal, say, that drug-testing is less than perfect, or that MoD is working on nerve gas? Should researchers publicise lines of research in the interests of science? Here, it might be fairly easy to weigh our commitment to the community against our obligation to protect our employer's interest even if

---

18 No doubt there are many who would question the premiss here, claiming that these technologies will never be profitable. Pursuing that would take us well away from the main line of argument I want to follow.

19 For the sake of argument, I am assuming that these concerns will make the technology more expensive, less saleable etc. I know that this need not necessarily be the case.

we have contractual obligations to maintain commercial secrecy.

With regard to invasive technologies, the problem is not quite of this kind. It is rather whether we should engage in research knowing what others are likely to make of it. Should we assume responsibility for what others do with our products? The trouble is that unlike work on gene splicing, say, or nuclear fission, the collective benefits of invasive technologies are undemonstrated.[20] However we do think that they will be saleable since some obvious benefits will accrue to some sections of the community (managers, supervisors etc.). What we seem to have become embroiled in is the relative priority to be given to different ethical imperatives, for business is ethically regulated too. Profits are not pursued in total disregard for any norms of conduct. It is just that these norms (e.g. 'caveat emptor!') may not necessarily sit easily alongside those adopted by the research community. Most of the time, the possibility of contradiction or tension is managed because the cases are so clear-cut. We think we know what to do in the nerve gas case. And because a faster algorithm for compressing video signals or a better way to store data does not appear to threaten the good of the community, we decide it does not have ethical implications. Invasive technologies seem to sit somewhere in the middle, which is possibly why Trojan Horse arguments are attractive when discussing them. They allow us satisfy both sets of demands at once.

Another version of the Trojan Horse argument goes like this. The negative potential of invasive technologies is recognised by all. Everyone fears the possibility of 1984 scenarios being applied to them, even the most hard-nosed, profit maximising businessman. This being so, these technologies offer a vehicle by which we as researchers could begin to play a major role in shaping business practices and priorities since it is we who have the expertise to fix how the technology presents itself. The attractiveness of the technology as product will make our skills highly prized. Part of the price we can extract might be a growth in our influence over the way the business conducts itself. Furthermore, since there is the possibility of tension between business ethics and scientific ethics, we should aim to reduce the gap between them. If we would prefer to change business ethics rather than those of science, one way to achieve this is to do

---

20 Indeed, some heretical souls argue this is true for the computational revolution more generally (Attewell unpublished)

research on invasive technologies. How far that leaves us in danger of succumbing to Eliot's last temptation is, though, unclear.

Is consequentialism a help?

So, after all this, am I any closer to answering my original questions? Well, yes and no. Clearly the utilitarianism which is appealed to as a method for choosing between ends is not much help in this case. We cannot envisage enough of the outcomes to be able to scale their benefits and so rationally choose between them. And without a utilitarian version of consequentialism, the original argument dissolves into special pleading, albeit of the Trojan Horse variety. This much is progress. Unfortunately, because we end up with a case of special pleading, we have now to confront and resolve the question of the relative priority which we as members of the scientific community working for commercial organisations are going to give to scientific and business ethics. And, furthermore, the ethics of business and science are not all of the radical individualist 'each to his own' kind. As I mentioned earlier, many businesses place now high value on customer satisfaction, quality processes, non-directive management and leadership, being a good citizen, and so forth. Of course, we can always be cynical and say that these are merely means to the old end of maximising profit. But why can't they be both means and ends? And what role can the commercial research organisations play in shaping business ethics and business practices to ensure a reconciliation between the demands being made?

This takes us away from the specifics of the case in hand and towards considering the character of the forms of legitimation themselves and what Max Weber (1947) called the difference between formal and substantive rationality. We look to formal rationality to legitimise the use of calculative methods as the basis of choice. We mathematise ethics, thus turning it into a technical discipline defined by technical expertise. A constant feature of the modern era, or so some would say, has been the attempt to transform ethical questions in technical ones, and then to look to formal rationality to provide means of answering them. We do this in two ways. We seek a sound set of principles (for instance the cost/benefit utilitarian version of consequentialism in the above discussion) which we can apply as a method and which will always provide the same answer independent of the questioner. Second, we look to formal reason embodied in other technical disciplines to remove our ethical dilemmas by transforming

their causes. I do not have to worry about switching off Granny's life-support machine if Granny can 'live' in suspended animation. I do not have to worry about the ethics of abortion if there is a 'morning-after' pill. Or do I? The extension of this approach to invasive technologies is to develop sophisticated methods and devices by which to give control to users. But does that really get us off the hook? Is tailorability and especially tailorability of access to and control over the functionality of invasive technologies really any better than the cynical use of quality processes and customer satisfaction? Are we sure it is not simply an ameliorating means to otherwise be unacceptable ends?

If we take a step or two backwards, it may be possible to get a broader view of what is going on here. The three clusters of considerations I have been mulling over, consequentialist ethics, the motivations for research and the morality of the marketplace all appear to appeal to the same set of principles — formal rationality. All seem to be predicated on the calculability of ends and the stability of a method, mathematical logic. Viewed from that point of view, the differences between them can appear to be frustrating since they ought, in principle, to be reconcilable. Given the common commitment to formal rationality, once there is an agreement on valuation (an agreement on the social arithmetic), the arguments should be over. And the methods of social arithmetic are themselves to be guaranteed by formal rationality — utilitarianism. The problem with this analysis is that it discounts the content in favour of the form. And it is the content which makes the differences. Using Weber's categories, what we have been discussing are the ways in which systems of *substantive* rationality jar against one another. Science, business and ethics are all (more or less) coherent and consistent systems: they are not the same system. Even if there is a process of homogenisation going on, that is only (a) partial and (b) largely restricted to vocabulary.

In a very clear way, then, these considerations give us reasons why we should want to take up the question of the ethics of research into invasive technology as an analytic or *critical* matter. As the two value systems I have just described increasingly jar against one another, the way this tension will be articulated is through consideration of the ethical issues involved. Indeed, one could say ethical dilemmas do not arise unless there is a clash or tension between values. It will no longer be enough to justify any technology solely on its own terms. That is, it will no longer be possible to treat a technology in isolation from the uses to which it is or can be put. An understanding

both of the breadth of the issues involved and of the rhetoric within which they are cast will enable us, as technologists, to keep our balance and sense of direction in the resulting maelstrom. From this point of view, the invasive technology question is just the first of what is likely to be an increasing number of issues which the research community may have to grapple with. This leads to a second set of reasons for considering ethical questions just now. As the number and level of ethical issues increases, it is more and more likely that matters of ethical choice will be imbricated with and come to influence questions which are wholly (or mostly) technical in nature. We need to be able to pick our way through the ethical issues in order to be able to determine, as a technical matter, just how far technical considerations should be shaped by them. The three most usual responses one hears (i.e. ‘not at all’, ‘it is too early to say’, and ‘in large measure’) are perhaps indicative of a view which treats the ethical questions as monolithic and beyond the purview of those wholly concerned with developing the technology. As we have seen, once one starts to reflect upon the arguments, that monolithic appearance soon dissolves. At that point, technical and ethical issues can be teased apart and examined both for what they are and for what they have to say to each other. When we have begun that process, we may well be in a position to begin to understand the complex intermingling of technical reasoning, reasoning in scientific research, in business, and more general frameworks of meaning or rationalisations loose in our society.

References

Attewell, P. 1991 Information, Technology and the Productivity Paradox. Talk given as a PARC Forum. Xerox Palo Alto Research Centre. Feb 21,

Clarke, R. 1989. Information, technology, and dataveillance. *Communications of ACM*, no 31, vol. 5, pp. 498-512.

Ellul, J. 1989. The search for Ethics in a technicist society. *Research in Philosophy and Technology*, vol. 9, pp. 23-36

Dunlop, C. and Kling, R. 1991. *Computerization and Controversy*. New York. Academic Press.

Fish, R., Kraut, R. and Chalfonte, B. 1990. The Video Window System in informal communication. *Proceedings of CSCW*, Los Angeles Oct 7 -10, ACM New York, pp. 1-12.

Gaver, W., Moran, T. Lövstrand, L., Dourish, P, Carter, K. and Buxton, W. 1991 Working together in media Space. Proceedings of HICOM, Jul 2, London.

Glover, J. 1986 'It makes no difference whether or not I do it'. *Applied Ethics* Peter Singer (ed.) Oxford. Oxford University Press.

Hare, R.M. 1981. *Moral Thinking*. Oxford. Oxford University Press

Harrison, S., Minneman, S., Stults, R., and Weber, K. 1991 VIDEO: a design medium. *SIGCHI Bulletin* vol. 21, no 3, pp. 86-90.

Lacayo, R. 1991. Nowhere to hide. *Time*. Nov 11.

Ladd, J. 1989 Computers and moral responsibility. In C. Gould (ed) *The Information Web: ethical and social implications of computer networking.* Boulder Colorado. Westview Press, pp. 207-227.

Mackay, W. 1991 Ethical issues in the use of video. Is it time to establish guidelines? In *Proceedings of CHI 1991* New Orleans, April 28 - May 2. ACM: New York, pp. 403-406

Mackie, J.L. 1977. *Ethics*. Harmondsworth. Penguin.

Mackintyre, A. 1990. Moral dilemmas. *Philosophy and Phenomenological Research*, vol. 11 Supplement Fall, pp. 876-882.

Mantei, M., Baecker, R., Sellen, A., Buxton, W., and Milligan, T. 1991. Experiences in the use of a media space. *Proceedings of CHI' 91*. New Orleans, May 3 - 5, pp. 203 -208.

Newman, W., Eldridge, M. and Lamming, M. 1991 PEPYS: generating autobiographies by automatic tracking. Proceedings of ECSCW, Amsterdam, Sept 26-27 .

Pullinger, D. 1989. Moral judgements in designing better systems. *Interacting with Computers*, vol. 1 no 1, pp. 93-104

Thompson, J. 1986. A defence of abortion. *Applied Ethics* Peter Singer (ed.) Oxford. Oxford University Press, pp. 36-57.

Toulmin, S. 1986 (originally 1956) *The Place of Reason in Ethics.* London. Chicago University Press.

Want, R. 1992 The Active Badge locator system. *ACM Transactions on Office HInformation System*. Forthcoming.

Weber, M. 1947. *The Theory of Social and Economic Organisation*. New York, Free Press.

Weiser, M. 1991. The computer for the 21st Century. *Scientific American.* Vol. 265, no. 3, pp. 94-105.