Virtual Reality

Virtual reality is neither frivolous nor necessarily the most important invention since transistors. (Pratt et al. 1995: 17)

In the computing world there has long been talk of the virtual. "Virtual machines" and "virtual memory" have been around for decades. These are essentially technical terms of little interest to anyone outside the computing profession. More recently we have begun to hear of virtual libraries, virtual conferences, virtual offices, and even virtual villages and communities. These have all resulted from the spread of the Internet and of socalled "cyberspace." People do not need to attend a conference physically. They can instead communicate via the Internet. Rather than commuting to a physical office, workers can telecommute, that is, work from home via a computer link to an "office." Libraries, instead of being collections of books, journals, and the like in buildings, can consist of electronically stored documents accessible remotely by computer.

These cases of the virtual do not really constitute virtual reality, although they do, to a greater or lesser extent, simulate parts of reality. A virtual machine much of the time gives the impression of being a stand-alone computer (in one use of "virtual machine"), and a virtual library performs the central functions of a "real" library, such as the storage of documents and methods for their retrieval. With virtual reality (VR) however, there is a more total involvement with the

computer-generated and -maintained environment. According to Pratt et al., VR is:

an application that lets users navigate and interact with a three-dimensional, computer-generated (and computer-maintained) environment in real time. This type of system has three major elements: interaction, 3D graphics, and immersion. (1995: 17)

They go on to explain this:

3D graphics, a form of computer output, let users "see" the virtual environment. Immersion refers to the user's feelings of "presence" in the virtual world. An immersive application convinces users that they are in a replicated environment.

This is pretty much a description of the current state of the art, concentrating as it does on the visual. In the future, however, more of the senses will probably become involved. By donning special equipment linking you via sensory interfaces to a computer, you may enter a world of VR. Depending on the program and in some cases on technological advances and refinements, you may then have the experience of climbing Mt. Everest, of making a century at Lord's, or of making love, without these things actually happening: you are not really climbing or cricketing or copulating, it just seems that you are.

What, if any, moral issues arise from the virtual in general, and from VR in particular? There is nothing of moral interest in virtual machines or memory. The next level, however, does raise some concerns, but ones considered elsewhere in this book. For example, intellectual ownership is a much discussed topic in relation to virtual (or digital) libraries, and we considered this in chapter 5. Again, there is the question of whether important human interaction is threatened by telecommuting, virtual libraries, shopping and banking over the Internet, and so on. This issue was also raised in chapters 8 and 9. The uses of VR itself range from education to entertainment, from science to sex. Its

scientific and educational uses are largely ethically uncontroversial, but some of its other uses are much less so. One of the best-known uses of VR is flight simulation, but there is an increasing number of others, for example, in surgery, in the treatment of the fear of heights, and in driving instruction (see Goble et al. 1995; Hodges et al. 1995; Kuhl et al. 1995). In all of these cases it would appear that behaviour in the real world can be enhanced by these virtual environments. Useful skills can be learnt without danger to the participants or to others in ways not previously possible, and this is surely uncontroversially good. At the other end of the spectrum

are things not so obviously good: virtual murder, rape and other violence, and other kinds of sex. Virtual violence, including murder and rape, may have some redeeming features. If it could be shown that people committing those virtual crimes were less likely to commit the real ones, then there would be some argument for those uses of VR. While society would undoubtedly be better off without any violence, presumably virtual violence is the lesser of two evils. It is an empirical matter whether virtual violence would increase or decrease real violence, and given that there is still no consensus on the effect of television violence, any strong evidence of the effect of virtual violence is some way off. Virtual sex is another matter, and it is on this that we will concentrate in this chapter for two reasons. First, it is controversial, and second, it highlights an important question: Is the real better than the virtual? Is an experience of the real better than an experience of the virtual? Whether VR should be employed to train pilots and car drivers is hardly a moral issue, nor is its use in, for example, giving the physically disabled the "experience" of hiking through the Andes. Enriching people's experience is one way to enhance the quality of their lives. Some uses of VR obviously have this potential, enabling valuable experiences that could not be obtained in any other way. But perhaps virtual sex is not in this category and instead lessens the quality of life. Perhaps again, all virtual experiences are like this.

Real virtual reality is not quite with us yet, but aspects of it are, and it will undoubtedly be developed further. But suppose that it is all with us now, and choices must be made and counsel given about its acceptable uses. Or suppose that it is about to be with us, and we have the opportunity to help to decide whether it will be or not; it is often lamented that certain technologies, notably in the biomedical area, have raced ahead of our moral, social, and political readiness for them, and so let us suppose that here we have the chance to be ready and to allow or disallow the technology to exist.

VIRTUAL SEX

VR authority Howard Rheingold invites us to imagine dressing for a hot night in the virtual village. You don a lightweight body-quit,

something like a body stocking, but with the kind of intimate snugness of a condom. Embedded in the inner surface of the suit, using a technology that does not yet exist, is an array of intelligent sensor-effectorsÑa mesh of tiny tactile detectors coupled to vibrators of varying degrees of hardness, hundreds of them per square inch, that can receive and transmit a realistic sense of tactile presence. (1991: 346)

They receive and transmit by exerting counterforces against your skin, counterforces corresponding to those you would encounter when handling actual objects. "You can run your cheek over (virtual) satin, and feel the difference when you encounter (virtual) flesh. Or you can gently squeeze something soft and pliable and feel it stiffen under your touch" (1991: 346). By plugging your whole sound-sighttouch telepresence system (which will include 3-dimensional

glasses and an appropriate auditory interface) into the telephone network, and "depending on what numbers you dial and which passwords you know and what you are willing to pay (or trade or do), you can find one partner, a dozen, a thousand, in various cyberspaces" (1991: 346); your physical body may be thousands of miles from your partner's, but you may touch and caress each other's virtual body.

You will whisper in your partner's ear, feel your partner's breath on your neck. You run your hand over your partner's clavicle, and 6000 miles away, an array of effectors are triggered, in just the right sequence, at just the right frequency, to convey the touch exactly the way you wish it to be conveyed. If you don't like the way the encounter is going, or someone requires your presence in physical reality, you can turn it all off by flicking a switch and taking off your virtual birthday suit. (1991: 346)

Rheingold goes on to say, however, and no doubt to the disappointment of some, that the required technology seems to be a long way off. Extensive fibreoptic networks will be required, though we will eventually be getting them anyway. The mesh of transducers, the tiny, high-speed, safe but powerful sensor-effectorsÑ"smart skin"Ñwill take some time to develop. But the biggest problem, he says, and the one that puts the above scenario into the early-to-mid-twenty-first century rather than the 1990s, is that extremely powerful computers will be needed to make the enormous number of calculations required to monitor and control hundreds of thousands of sensors and effectors. "Every nook and protuberance, every plane and valley and knob of your body's surface, will require its own processor" (1991: 347). Suitably sophisticated olfactory and gustatory sensors and effectors, to complete the range of sensory communication, would presumably be technically challenging too. We shall just have to be patient.

Note that in Rheingold's scenario one is paying or trading or doing something in exchange for the experience; probably, in fact, one would be paying. So we have not only virtual sex, but virtual prostitution. And certain further concerns might arise on the strength of this. Before we tackle those, let us take Rheingold~s scenario one large step further. Imagine that

the virtual body you are virtually caressing is not a representation of an actual person who is virtually in contact with a similar representation of your actual body, but a virtual body behaving as it does merely as a result of a very sophisticated computer program, written or chosen presumably according to your wishes. There is no person on the other end, or at least no person having the sensation of contact with your body corresponding to your sensation of contact with hers or his. We have no idea when such a program and thus such an event might come about, but they seem imaginable. And if and when they do, they will provide virtual sex of a different kind: Your virtual partner will be animated not in a symmetrical way by someone who is doing essentially what you are doing, but in an asymmetrical way by the programmers and technicians who are not doing what you are doing at all. Let us distinguish these two kinds of virtual sex by calling them SV (for symmetrically virtual) and AV (for asymmetrically virtual) sex. SV sex so defined will not of course necessarily be symmetrical in any other way, as the partners may differ markedly in terms of motivation and enjoyment, just as in actual sex. In AV sex, the virtual partner may be more or less closely modelled (perhaps according to the client's request) on some real person, with any imaginable positive or negative degree of that person's approval or consent. But however close the modelling and however consenting the original person may be, it is appropriate to classify such a case as AV rather than SV, for the original is not having a reciprocal sexual experience, and may not even be alive at the time. Where the modelling is close and there is no

consent, some may see it as a special case of rape; others may see it as morally distinguishable, without necessarily defending it.

Virtual Sex and Pornography

Virtual sex obviously invites comparisons with pornography. Various types of AV sex may be seen as respectively corresponding to various types of pornography: We may distinguish, for example, fiction or drawings of imaginary people on the one hand from drawings or paintings from life, photographs, and films on the other. And of course in the latter cases models may have been somehow exploited or otherwise harmed. Whatever the circumstances of its production, however, one of the standard objections to pornography is that it causes or may help to cause the consumer to commit sexual assaults, perhaps of the very kind that he (usually he) has enjoyed pornographically. Should there be a similar concern about either SV or AV sex? It is often pointed out that on such tricky matters we must be careful in drawing conclusions as to what causes what. Even if sexual offenders tend to be consumers of pomography, this does not show that the latter contributes causally to the offences even if the offenders say and believe that it does: The interest in pomography and

the subsequent criminal behaviour may both be caused by some underlying personality factor absent in the many consumers who do not commit offences. Maybe so; but the would-be censor or opponent of pornography usually replies that it may at least act as a trigger, having an important causal role despite not being the most fundamental factor. And when virtual sex is available, may it not do so as well?

On the other hand, it is often claimed that pornography may actually serve as a "safety valve," by diverting or satisfying the lust of some potential rapists and molesters; it might be suggested that if virtual sex of either kind is realistic enough, it could do this job quite effectively. It may also seem that any virtual sex acting as such a safety valve is likely to be AV rather than SV, since not many women would wish to be "virtually raped" and not many children would wish to be, or would be allowed to be, "virtually molested." But some women may be prepared to take part in virtual rape, especially for money, just as some prostitutes presumably go along with some clients' rape fantasies and some actresses and models take part in simulated rape scenes; and some adults might, again especially for money, allow their virtual selves to be metamorphosed by VR technology into those of molestable children. Such SV safety valves may look very squalid, but they may be safety valves nonetheless. We should also consider the possibility that, to serve as a safety valve preventing a rape, an episode of either SV or AV sex need not be one of virtual rape, in at least some cases virtual consensual sex may do the trick. But this idea invites the reminder that many rapes, perhaps most of them, appear to be motivated primarily not by a desire for a sexual experience nor even by a desire to have actual sex, but by a desire to exercise power over an actual woman, power that in the rapist's view is most satisfactorily exercised in a sexual manner with or without additional violence. Where something like this is the rapist's motivation and he wants to exercise this power rather than merely experience the feeling of exercising it, then of course no VR experience known to be such will do. Another potential offender who will not be diverted by a VR safety valve is the one who, having experienced virtual sex (perhaps virtual rape), wishes to try out the real thing to see whether and how the two experiences differ, and is prepared to rape accordingly. As in the previous case, a high degree of sophistication and realism in the VR experience will not divert such a person, and here it may only encourage him. But of course he would have to be not only curious in this matter, but prepared to rape or otherwise assault someone to satisfy that curiosity. And it would seem difficult to be sure that the number

of offences thus caused by the availability of VR sex would exceed the number prevented by its operation as a safety valve.

Someone who opposes indulgence in virtual sex or would wish to ban it may perhaps be confident of having good grounds for the belief that its bad consequences will outweigh its good ones. The bad ones may be

thought to include not only criminal behaviour, but the treating of women and maybe children as sex objects, which pornography is often thought to help to cause. A hedonistic or Benthamite utilitarian, of course, would have to count the pleasures of the consumers among the good consequences; a wider view might count the sexual exploitation in some forms of AV sex as a bad thing, even where it causes no distress. Alternatively, one might hold that a significant chartce of harm is enough to justify opposition or prohibition, even if this chance falls short of an even money bet. Such a position is vulnerable to a countermove in terms of a significant chance of averting harm by way of a safety valve, and also to the question of whether it can be established that the chance of harm is significant enough to justify the restriction of personal liberties, especially those of people whose other behaviour and personality give no cause for concern. But there is a different type of position that the opponent of virtual sex may take, along with or instead of either of the two just consideredÑthe position that, quite apart from their consequences, virtual sexual experiences are in some way objectionable. How so? Two compatible answers may be offered.

First, it may be thought that such experiencesÑlike various other sexual experiences including the enjoyment of pornographyÑare bad, and that it is wrong or sinful to seek them, because they contravene the purpose of sex as discoverable by biology or divine revelation or the "natural law." Sex is for procreation, it is said, and it is wrong to seek sexual experiences that do not at least leave that possibility open. But there are certain well-known difficulties here. Biology tells us only what actually happens in nature, not what we ought or ought not to do; divine revelation is hard to verify and (even in the most respected texts) not very explicit anyway on such matters as this, so that it requires interpretation by supposed authorities; and natural law, however understood, is not very obviously able to steer between one of these two problems and the other.

The Real and the Right

Alternatively, one may see virtual sex as a bad thing because it is essentially fake. The objection here does not specifically concern its sexual nature, but the fact that one is accepting or seeking a fake thing instead of a real one, a shadow instead of a substance, even if the shadow seems by courtesy of VR to be mighty substantial.

Several things may be said in response to this. First, in what sense is SV sex objectionably fake? Is the fact that the partners' sexual encounter is not directly body-tobody decisive, despite the fact that they are behaving symmetrically in our important sense, with each partner's sexual sensations being intentionally caused by appropriate behaviour of the other? AV sex may be a different matter, but it is not clear that SV

sexual intercourse is any more fake intercourse than telephone conversation is fake conversation because the sound waves the talkers emit do not directly strike one another's eardrums, but rather are translated into and then out of electrical signals. But let us concede for the sake of argument that SV sex as well as AV sex should be seen as fake. Let us also concede that, other things being equal, a real X is better thanÑ more worth having thanÑa fake X: The real

Mona Lisa is more worth having than a forged Mona Lisa, real teeth are more worth having than false teeth, and so on. But what if the real thing is not available? If your real teeth have gone, you may well settle for a set of false ones, provided that the appearance and comfort are up to scratch, and nobody would think any less of you for that. What about the Mona Lisa? Here one might think that, rather than a fake that implausibly purports to be the original, it would be better to have what is obviously just a print; but a lot depends here on one's interests. A forgery might be of special charm or interest (or monetary value) because of its quality or who did it or owned it, while a print might be better to have for certain other purposes such as display in the school art room. Must the original, however, be preferable to either? Even if one could acquire the original by lifting it from the Louvre, one might very much prefer not to do so for various reasons: its safety, fear of arrest, the interests of wouldbe viewers, and so on.

Here, then, is a second important point: that while the real thing may be better to have than a fake or copy where other things are equal, other things may very well not be equal, as various considerations may affect one's overall preference. With respect to certain types of things, and in certain circumstances or for certain purposes, one may reasonably prefer to have a fake or copy of some kind rather than the real thing. One may choose to look at photos or films of exotic places rather than incur the trouble and expense (and maybe danger) of going there oneself; one may wear imitation fur for financial or moral reasons; one may eat imitation butter (margarine) for moral or health reasons. Why should someone not be prepared to make a similar choice, at least sometimes, as to sexual experiences? It is often noted that virtual sex, like currently available entertainments such as anonymous "fantasy phone calls," would be free of the risks of serious unwanted consequences such as pregnancy and disease. So some might be attracted to a form of sexual expression that, while

ferior "in itself" to some more or less achievable alternative, is on balance preferable at least some of the time because it does not expose oneself or a partner to such risks. Particularly (though not necessarily) where one has the partner's interests in mind, one's preference here may actually be a moral one.

Quite apart from this there is a third important point to be noted. Even if any reasonable person would judge a certain choice to be inferior to

some other, this does not mean that it would be morally wrong to make the former choice either occasionally or regularly. Eating at a fast-food restaurant may be thought inferior to dining in style at a slow-food one, but this does not mean that it is wrong to choose the former where one can afford the latter. We may be sad that someone makes the worse choice on some occasion or on most, and we may feel entitled or even obliged to say so. But this sentiment may well fall short of a judgement of wrongness; your regret at your friend's choice of partner or pastime may not amount to the feeling of blame or resentment, however mild, that is characteristic of judging something morally wrong. It would be a very demanding morality that extended judgements of obligation and wrongness to cover all such cases. And might one not see virtual sexÑof either the SV or the AV kindÑas a rather sad thing for a person to choose where the real thing is available, without seeing the choice as morally wrong?

Fourth, even if we do judge it morally wrong, this does not mean that we must favour any coercion or legal sanction. The proper province of the law may be seen as stopping short of prohibiting immorality as such: One may adhere to J. S. Mill's celebrated doctrine, which we came across in an earlier chapter, that we should have the liberty to do as we like "without impediment from our

fellow-creatures, so long as what we do does not harm them, even though they should think our conduct foolish, perverse, or wrong." Unless it is argued that virtual sex causes or threatens to cause harm of some kind, such as mentioned earlier, to prohibit it would be an exercise of what is now called legal moralism, condemned by Mill and many others in the liberal tradition. The wrongness of a person's behaviour, says Mill, is, like folly, a good reason "for remonstrating with him, or reasoning with him, or persuading him, or entreating him, but not for compelling him, or visiting him with any evil in case he do otherwise" (1859: 11). The line between some forms of persuasion and some ways of "visiting with evil" may be rather fine, but Mill would certainly wish society to attempt to draw it in the case of behaviour that, even if widely thought to be wrong, cannot be plausibly seen as harmful to others. Various sexual activities such as homosexuality and promiscuity have been suggested as examples; virtual sex might be offered by some as another.

Virtual Prostitution

Both SV and AV sex will have their commercial manifestations. There will be SV prostitution, in which the real person on the other end may be regarded as akin to a prostitute, giving and having a sexual encounter for money. AV sex will presumably be typically commercial, and the technicians and marketers who arrange it will be akin to pimps. Cases will

include those in which the virtual partner is modelled on a fully consenting person who, if being paid, will thus have some similarity to the SV prostitute, though one step further removed from the client; those in which the virtual partner is modelled on a nonconsenting person who would understandably see herself (usually herself) as being distastefully exploited; and those in which the client's virtual partner is not modelled on anyone in particular, so that there is a virtual prostitute but no real person who is virtually a prostitute.

Real prostitution, the sort familiar to human society as its oldest profession, is morally condemned on various grounds. As in the case of virtual sex, the objections may be divided into two categories. On the one hand, it is often claimed that prostitution is in some way harmful to one or more of the parties directly or indirectly involved. Alternatively, or additionally, it may be held that it is somehow a bad thing in itself, quite apart from any harm it may do to anyone. Let us briefly consider these types of objection in turn and their bearing on virtual prostitution.

An important point often made is that much of the harm allegedly caused by prostitution may be caused by factors involved in only some types or instances of it, or by its being illegal or socially frowned on, not by acts of prostitution as such. No doubt some prostitutes are exploited in being coerced, or in pimps and clients taking advantage of their inability to keep themselves and their children in any other way. But could it be said that even the up-market prostitute who prefers that occupation to a more modestly paid but available altemative is being exploited, in a sense substantial enough to mean that she is being wronged? To say that she is however strong her sense of free choice may be, is to suggest that no woman could rationally and autonomously make that choice. There seems to be no good reason to accept that suggestion, and good reason for many prostitutes to be offended by it. The choice may not be one that w

ourselves would ever care to make, but the same goes for many a choice that we would not dream of calling irrational.

SV prostitution could certainly be exploitative in very unpleasant ways,

although, given the hi-tech and possibly up-market nature of the exercise, one may guess that the worst types of case would be found les

commonly there than in the real thing. And it could be just as nonexploitative as the real thing too. AV cases of the kind in which the virtual partner represents a recognisable, actual person could be very exploitative, either because that actual person has been subjected to unfair influence to be the model, or (perhaps more probably) because she has not consented at all and the virtual prostitution consists in her virtual services being sold, not in her selling them. But where she has given her fre

and informed consent, then she too (like her SV and actual counterparts) can hardly be seen as exploited. And in AV cases of the more "fictional"

kind, in which the partner is not modelled on anyone in particular, there is nobody to be exploited.

According to moral objections of the second type, prostitution is simply wrong in itself. Why? It is hard to see why the selling of sex should be as such immoral, and, as often observed, hard to see how it differs morally from trading sex for other goods such as the material support and social status of a marriage, which some spouses more or less frankly do. Perhaps the wrongness of commercial sex, as well as much other sex, is seen as due to its lovelessness. It is not clear that it is necessarily loveless, at least on the part of the client, but let us agree that it is at least characteristically so. Well, what is wrong with loveless sex? Some have argued that just as sexless love may be perfectly all right, so may the converse: Love is a wonderful thing and so is sex and it is fine when the two coincide, but nothing is amiss when they do not unless someone is or may be hurt or betrayed, in which case the objection is not to the separation of love and sex per se but to the hurt or betrayal. It may well be that the idea of an entirely loveless sex life has little appeal for sensitive and loving people, but this does not mean that occasional loveless sex is unappealing to them too. Even if it is, more than the knowledge of this attitude would seem to be required to justify the claim that loveless sex is inferior. And there are several further points to be made that parallel those made earlier concerning the fake and the real. That which is inferior where other things are equal is not necessarily to be avoided where other things are not equal: What if embarking on a loving sexual relationship would create risks and harms for the other party? One might be about to go overseas for years, or off to war or a monastery. And of course in a given situation a loving relationship may simply not be in prospect. Next, to choose the inferior is not necessarily wrong; we may regret that someone chooses occasionally or exclusively to have loveless sex with prostitutes or casual pickups rather than developing a relationship in which the sex will be loving, but this is not the same as thinking it wrong. And finally, of course, even if we do think it wrong to have sex lovelessly and commercially, this is not to believe that it should be illegal. One might see prostitution as wrong but in essence harmless, and thus in that realm of private morality that is no business of the law. One might also see positive benefits in toleration, including the reduction of police corruption and other associated crime, and, as many have claimed, the protection of "honest women" from much molestation. One of us once met at a soiree a pleasant elderly lady who, commenting on a recent spate of sex crimes in a large city, said she thought there ought to be more brothels. "And they should befree!"

The day of the free brothel may be some way off; that of the virtual brothel may be a little closer. Perhaps the cost of the technology will even

tually drop so much that the AV brothel is practically free anyway. Its commodity may fall somewhat short of an ideal many of us have about sexuality at its best, and in fact may be seen as multiply removed from that ideal in being fake, commercial, and loveless; but more than that is needed to justify moral condemnation, and more still to justify prohibition of its operation or of the development of its sensational technology.

VIRTUAL REALITYANDTHE QUALITY OF LIFE

VR in general has the potential to enhance the quality of life, particularly in education and science; that is almost certain. Our argument about virtual sex suggests that this use of the technology too could make life more rewarding at least for some. Even if virtual sex is not as good and satisfying as the real thing in a loving relationship, it could nevertheless enhance the quality of life for those for whom a loving relationship is beyond reach, for whatever reason. It also has the potential for making life safer, especially for women and children, by providing sexual outlets that harm nobody. While the world would be a better place in all sorts of ways if everybody were in a loving relationship, this ideal world does not exist and almost certainly never will. We need therefore to use technology in ways that will make it a better place, even if not the best one. Although this chapter has concentrated on virtual sex and concluded that it does have a place, that is not really the main conclusion of the argument. The main conclusion is that, in general, there is nothing wrong with simulated or virtual environments. Virtual sex may not be as good (in some sense) as real sex, or virtual travel as real travel, but nothing follows from that about their rightness or wrongness. They may be a lot better than missing the experiences altogether. There are dangers of course. It would seem undesirable that someone should live his or her whole life in virtual reality if the real thing is available. Why it seems undesirable is not so easy to explain. Perhaps it is because living like this is parasitic on the rest of society. Another reason is that perhaps it is a denial of humanity. Such a person would go through life without facing and overcoming any challenges. But that is another issue. Our concern here has been to argue that there is nothing wrong with VR in itself, even for use in very personal areas.

SUMMARY

The discussion of VR has centred on virtual sex because it is controversial, and if VR can be defended here, it can probably be defended in almost any area. We argued that while it (virtual sex) might be poorer than the real thing, that in itself does not mean that it has no value. We found no arguments to support banning it, and it was suggested that it does

have the potential to improve the quality of life for some. These argu

ments can be generalised to most applications of VR. It can improve our

living, working, and playing environments.

REFERENCES

Goble, John C., Ken Hinckley, Randy Pausch, John W. Snell, and Neal F. Kassell. 1995. Two-handed spatial interface tools for neurosurgical planning. Computer 28 (July): 20-26.

Hodges, Larry F., Rob Kooper, Thomas C. Meyer, Barbara O. Rothbaum, Dan Opdyke, Johannes J. de Graaff, James S. Williford, and Max M. North. 1995. Virtual environments for treating the fear of heights. Computer 28 (July): 27-34.

Kuhl, Jon, Douglas Evans, Yiannis Papelis, Richard Romano, and Ginger Watson. 1995. The Iowa simulator: An immersive research environment. Computer 28 (July): 42-48.

Mill, John Stuart. 1859. On Liberty; page citations to edition of David Spitz, 1975.

Pratt, David R., Michael Zyda, and Kristen Kelleher. 1995. Virtual reality: In the mind of the beholder. Computer 28 (July): 17-18.

Rheingold, Howard. 1991. Virtual Reality. London: Secker & Warburg.

FURTHER READING

Antoff, Michael. 1993. Living in a virtual world. Popular Science (June): 124.

Communications of the ACM 38 (November) 1995. Special issue on digital libraries.

Communications of the ACM 39 (May) 1996. Special section on Virtual Reality.

Corliss, Richard. 1993. Virtual man! Time (November 1): 80-83.

Dvorak, John C. 1992. America, are you ready for simulated sex and virtual reality? PC Computing (May): 78.

Larijani, L. Casey. 1994. The Virtual Reality Primer. New York: McGraw Hill.

Pimentel, Ken, and Kevin Texeira. 1993. Virtual Reality: Through the New Looking Glass. New York: McGraw Hill.

Spring, Michael B. 1995. The virtual library. Fantastic Futures: Virtual and in the Flesh. National reference and information services conference, September 13-15. Presented by the Reference and Information Services Section of ALIA (RAISS).

Wexelblat, Alan, ed. 1993. Virtual Reality: Applications and Explorations. Boston: Academic Press Professional.

Minds, Machines, and Morality

When a man Reasoneth, tree does nothing else but conceive a summe totall, from Addition of parcels: or conceive a Remainder, from Substradion of one summe from another.... Out of all which we may define, (that is to say determine,) what that is, which is meant by this word Reason, when wee reckon it amongst the Faculties of the mind. For REASON, in this sense, is nothing but Reckoning (that is, Adding and Subtracting) of the Consequences of generall names agreed upon, for the marking and signiping of our thoughts. (Hobbes 1651: 1 1~1 1)

The issues discussed in this book cluster around the idea of a computer as an information processing machine. We have looked at questions related to the processing and communication of information, and at some related to the information generated by computers. The environment, both real and virtual, created by various kinds of computer use was also examined. It remains now to consider the nature of the machines themselves. The use of these information processing machines raises a host of moral questions, some of which we have

seen. Do the machines themselves raise any? Can they, in principle, be developed to a stage at which they should be treated as moral agents? This is one of those issues, mentioned at the beginning of the book, that has not arisen as a practical problem yet, but in all probability will, if there is continued research and development in intelligent machines. It is our contention that likely problems should be discussed, so that if or when they surface, we have some idea of how to deal with them. The central question here is whether

or not computers can have minds as humans do, and it is to this that we now turn.

In previous chapters we looked at some ethical questions about computers performing tasks normally done by humans or that may indeed go well beyond the abilities of even the smartest humans. They perform feats of memory, of calculation, of deduction, and so on, that for speed and difficulty leave human minds far behind. Could computers themselves have minds? Could they reason as we do, going through a sequence of thought? Could they have any thoughts at all? What about feelings or emotions, such as irritation or anger or envy or love? When a system error occurs in my computer and it says "Sorry," it is behaving as if it is sorry. Nobody believes that it actually is, but could there be a computer that not only says "Sorry" but is sorry?

We shall approach this question by way of considering the nature of mentality in a human being. Each of us has this dimension of life, in that we think, feel, believe, and so on. What is the nature of this mentality? And what is its relationship to our physical body, and in particular to the brain, which clearly has a more intimate connection with our mentality than does the liver or the thorax? Many answers have been given to these questions, many different theories advanced on the "mind-body problem," as it is often called. We shall survey a few of these theories, inquiring in each case as to the compatibility of the theory with the idea of machine mentality.

THE MIND-BODY PROBLEM

Dualism

According to a very popular view, a human being has two fundamentally different components. One is the body, a physical thing with size, shape, mass, and so on, and including many physical parts and organs. The other is the mind, regarded as a nonphysical, immaterial thing, a thing without size or shape or mass but a thing nonetheless. The mind (or soul, as dualists sometimes call it) is the bearer or arena of our mentality of our consciousness. Many dualists willingly concede that the mind and the body are closely linked, especially via the latter's brain; what goes on in the brain may have quite an effect on the mind, and vice versa. An injury to the brain, for example, or an encounter with alcohol, may produce odd perceptions in the mind, while a mental event such as a decision to go home will normally cause electrochemical events in the brain that in turn cause messages to go out to the muscles and the body to move in the direction of home or the means of getting there. But despite this causal interaction, intimate though it is, the mind and the body are different things. Many dualists hold further that they are not of equal importance in the

constitution of a personÑthat your mind in fact constitutes the essential you. One of the most famous dualist philosophers, Rene Descartes, said in his Sixth Meditation, "I am really distinct from my body, and could exist without it" (Descartes 1642: 115). And like many other dualists, Descartes held that the mind (i.e., the person) does survive the body, and indeed survive it for ever. Dualism is consistent with many religious beliefs, and this helps to account for

its popularity, though it also appeals to those who hope for survival of death in some form without any specifically religious focus. (Some people hope for life after death just so that they may see their loved ones again.) Yet it is important to be aware that the distinctness of mind and body is one question, the survival of the latter by the former is another, and survival forever is another one again. A bolt and its nut, though intimately related, are two distinct things that could survive apart; but it does not follow that either will survive forever, and they may even perish together by rusting away or melting. Nevertheless, dualism is obviously congenial to the ideas of life after death and immortality, and certainly more so than several other views we shall be mentioning.

Despite its appeal, dualism faces many problems for the reasons enumerated above and others. Essentially, a lot of the difficulties come down to one of the following: First, there is the problem of what sort of thing an immaterial mind or soul could be. It is tempting to try to define it negatively, by saying that it is invisible, intangible, and so on, but this fails to distinguish the mind from nothing at all, because nothing is also invisible, intangible, and so on. Positive descriptions of the mind, especially descriptions of what it is as distinct from what it does, are not easy to come by. Second, there is the problem of its relationship with the body, a relationship admitted by most dualists to be very close despite their utterly different natures. On the latter in particular Descartes was pressed by Princess Elizabeth of Bohemia, with whom he corresponded in later life, and was driven to say to her in exasperation that the union of soul and body is perfectly well understood until we come to do philosophy, and so perhaps we'd better not do too much philosophyÑa rather startling suggestion from a man often called the Father of Modern Philosophy.

The difficulty of establishing the nature of an immaterial mind does not show that there is no such thing, and the difficulty of describing its relationship with the body does not show that there is none. If dualism of some sort is the correct view of human mentality, where does the computer stand? Surely, many say, a computer could not have an immaterial soul. How could the engineers at IBM or Apple possibly make and install one? Presumably they could not. But does it follow that a computer could not acquire a soul in some other way? When you mix concrete, you do not add the hardness as one of the ingredients like the cement and the sand and the water. The hardness comes about as a result of your mixing those three things together. Now admittedly the hardness is not a thing in the

way that the dualist's mind or soul is supposed to be, but rather a quality, and a physical quality at that. Nevertheless, the example shows that in assembling X you can give it Y without doing so directly: Y may be a product of what you are doing directly. It may even be an unintended product, depending on the case. So, even though a computer builder could hardly make or take a soul and put it into a computer, could not a suitably sophisticated machine produce one for itself?

Many dualists will respond to this by saying that this is impossible, and that human bodies do not produce souls for themselves either. The soul is made by God and implanted by him at conception, according to the Catholic Church, at least since a ruling by Pope Pius IX in 1869; prior to that, the church accepted the ancient idea that "ensoulment" took place on the fortieth day of the pregnancy in the case of a male foetus and on the eightieth day in the case of a female one. But whenever it actually occurs, ensoulment is done by God, and he would never do it for anything nonhuman. He would certainly never do it for a machine. But why is this out of the question? How do we know what God might or might not choose to do? The soul, if it exists, is such a mysterious thing

anyway, as is its relationship to the living body, as also are God's soul-making and soul-installing activities, that it seems highly presumptuous on our part to think we can be sure that he would never give one to a nonliving thing. Perhaps it would have to be sufficiently like us in some key respect, but who can know what that respect is or what degree of similarity would be sufficient?

Behaviourism

Let us now consider another significant view of mentality and its implications for computer mentality. Dualism is attractive to many people especially those with a religious belief or an interest in an afterlifeÑbut unattractive to others, who reject the idea of the body somehow containing and being directed by an immaterial mental substance, a "ghost in the machine." Among those who reject dualism are many psychologists of an empiricist temperÑdealing only with, and believing only in, that which can be publicly observed, measured, and so on. The resident psychologist on a tabloid newspaper was once asked by a reader why he never said anything about the soul. He replied that he had looked up "soul" in his dictionary, and it said "the immaterial part of man." He then looked up "immaterial," and it said "irrelevant," so we need not bother with it. His reply traded of course on the ambiguity of "immaterial," but he may have been well aware of that.

Rejecting dualism, many psychologists and philosophers have embraced in one form or another a view known as behaviourism. Essentially, the behaviourist position is that there are not two things, the mind and the body, but only one, the body. Thanks largely to its brain but also of course to its nerves and muscles and bones and the rest, it behaves in all sorts of waysÑwalking, talking, leaping, weeping, and so on. Our mentality, says the behaviourist, is not something behind this bodily behaviour, a mysterious inner life that somehow directs it; rather, it just is our bodily behaviour in some way. To believe that today is Tuesday, for example, is to behave in a Tuesdayish way, which will involve responding appropriately to questions and doing other things one does only on Tuesdays; to be amused is to engage in amusement behaviour, such as laughing or smiling. The bodily behaviour is no doubt caused by neural and muscular events, with the brain playing an important role; but whatever its causes may be, mentality itself is constituted by the outward behaviour.

While dualism is commonly thought (perhaps wrongly, as we have seen) to exclude computer mentality, behaviourism seems to allow it, at least in principle. If a computer, or at least a computerised robot, could walk, talk, smile, and so on, then it would be displaying the behaviour of a being with mentality and so it would have to count, on the behaviourist view, as having mentality. Of course someone might say that a robot's smile is not a real smile, for a real one has to be on a human face and not on a metal or plastic one. But this would be to settle the issue against the robot simply by definition; it is not going to be allowed to qualify by displaying the appropriate behaviour, because the appropriate behaviour is by definition human. And a point scored against the robot in such a way is no point at all. We might just as well say that a robot could never be said to play the piano, no matter how beautiful the rendition of the Moonlight Sonata it produces by striking the keys with its metal fingers, because playing the piano is by definition an exclusively human activity.

So behaviourism, unprejudicially understood, seems to also allow the possibility in principle of machine mentality. How possible it is in practice

will depend not only on technological developments but also on the level of sophistication in the behaviour required. Unfortunately, however, behaviourism seems very easy to refute. A mental state (such as feeling or believing something) cannot be in any way identical with the appropriate physical behaviour for the simple reason that it is quite possible to have either one without the otherÑa mental state without the physical behaviour in which it is normally expressed, or vice versa. Take, for example, amusement. This state may often or characteristically be manifested in certain outward behaviour, such as smiling or laughing. But you might feel amused without showing it at allÑsomeone has committed a terrible faux pas, and tact or propriety require you to pretend not to have noticed. Or you might display amusement behaviour without being at all amusedÑyour wealthy prospective father-in-law has just made a very feeble joke, and you fall about laughing. Another possibility is that you are acting in a play, and you have to laugh at something that amused you

the first time you heard it but not any longer. The state of amusement then is something different from the amusement behaviour; they often occur together and are causally related, but in various circumstances either one may occur without the other. The same goes for other mental statesÑfeelings, emotions, thoughts, beliefs. They are connected with our behaviour but distinct from it, and have an inner status.

A similar point may be made concerning computers. Outward behaviour is one thing, mentality is another. In his discussion of computer mentality, John Searle offers his now celebrated parable of the Chinese Room. Imagine, he says, that you are locked in a room with several baskets full of Chinese symbols (perhaps on pieces of paper). Further collections of symbols are passed into the room from time to time. You understand neither the symbols in the baskets nor the ones coming in, but you do have a rule book (in English) which gives you precise directions for responding to the incoming symbols by assembling ones from the baskets in certain ways and passing them out again. Unknown to you, the people outside the room call the collections they are sending in "questions" and the ones coming out from you "answers." If the rule book is good enough and you become quick enough, the answers you send out may be indistinguishable from those of a native Chinese speakerÑbut, says Searle, you are not learning any Ch~nese.

Following such rules then, no matter how well you do it, is not enough to give you an understanding of Chinese. You may acquire an excellent knowledge of Chinese syntaxÑthe formal rules for arranging the symbolsÑbut you have no semantics, for you do not know what the symbols mean. And when a computer follows the set of rules that constitute its program, then no matter how sophisticated that program and how excellent the machine's execution of it, this is not the same thing as having an understanding of what is going on, and thus not the same thing as having a mind. For, says Searle,

There is more to having a mind than having formal or syntactical processes. Our internal mental states, by definition, have certain sorts of contents. If I am thinking about Kansas City or wishing that I had a cold beer to drink or wondering if there will be a fall in interest rates, in each case my mental state has a certain mental content in addition to whatever formal features it might have. That is, even if my thoughts occur to me in strings of symbols, there must be more to the thought than the abstract strings, because strings by themselves can't have any meaning. If my thoughts are to be about anything, then the strings must have a meaning which makes the thoughts about those things. In a word, the mind has more than a syntax, it has a semantics. (1984: 31)

And so, Searle concludes, "no computer program by itself is sufficient to give a

system a mind. Programs, in short, are not minds, and they are not by themselves aufficient for having minds" (1984: 39).

Mind-Brain Materialise..

If mentality is something innerÑsomething causally connected with outward behaviour but distinct from itÑthen is the dualist view correct after all? Not necessarily. The dualist view is that the mind is an inner thing of an immaterial nature. But could it not be something of a material nature, namely the brain? According to a third view of the mind-body relationship, mental states such as belief or amusement are simply states of the brain, and mental processes such as inferring or speculating are simply processes in the brain. Mental goings-on are not just correlated with or connected with physical goings-on in the brain, as most dualists acceptÑ rather, they are those brain goings-on. Such a view is sometimes called materialism, though this term is also used in a wider sense to cover any view to the effect that the only constituents of a human being are material ones. Behaviourism as defined above would count as a variety of materialism in this wider sense. The view we are now considering is certainly also materialistic in this wider sense, but differs from behaviourism in identifying our mentality with states and processes of the brain rather than with aspects or segments of outward behaviour. So we shall call it mind-brain materialism. Like other forms of materialism, it seems to rule out the possibility of life after death because the finish of your brain would be the finish of your mentality and of your personality and of you yourselfÑunless of course you see your survival in the continuance of the elementary particles that currently make up your body but will eventually make up other things. What they make up will depend on whether your body is burnt or devoured by worms or crocodiles or whatever. But our concern here is more with two other questions: (1) Does mind-brain materialism allow the possibility of machine mentality? (2) Is mind-brain

materlallsm allow ¥ll~ 1UU;~:~lulllty wl AlI~

materialism true?

To the first question, mind-brain materialists may not all give the same answer. Mental processes happen to be, on their view, brain processes; could they also occur in something that is not a brain but aufficiently like one? How similar would it have to be? Purifying blood is a process that takes place in the kidneys, but it may also occur in something that is not a kidney but relevantly and aufficiently similar to one - a dialysis machine; it does not occur in something which is like a kidney in other ways, such as a plastic kidney in one of those life-sized cutaway model bodies used in anatomy classes. Mental processes do not (presumably) occur in the model body's model brain. Could they occur in something sufficiently like a real brain, in whatever respects are relevant? And what are they?

Some will say that no computer could possibly be like a real brain in the relevant respects, for a very important one of those respects is that the latter is a living organ. Blood might be purified by a machine, either outside or inside the body, but mentality is something altogether different. It simply could not occur in anything that is not alive. Why not? With respect to dualism we wondered why a computer might not produce an immaterial soul or be given one by God; in the present case we might similarly wonder why a computer could not equal a brain in having mental processes. The supposed identity between mental and neural processes is a difficult thing to comprehend even in a human, and the defender of the possibility of machine mentality may ask why a similar identity. could not hold between a mental process and an electronic one in a suitably

sophisticated computer.

Now for our second question: Is mind-brain materialism actually true? It has the attraction for many people of giving a simpler account of human nature than the dualist's, while recognising (as the behaviourist does not) that there is an important difference between being in a certain mental state and behaving outwardly in the appropriate way. It also offers a neat explanation, which the dualist can hardly do, of the intimate connection between the brain and the mind. They are not two utterly different thing

one with size and shape and mass and the other without, which are nevertheless mysteriously bound up with one another, "intermingled," as Descartes had to agree. Rather, they are one and the same thing. So

what similarly, what better explanation could there be of the intimate connection between the lives of Clark Kent and Superman than that they are one and the same person?

Despite its attractions, this version of materialism comes up against some important objections, of which we shall raise two. First, there is what is often called the Privacy Objection. It is said that there is a fundamental difference between what goes on in the brain and what goes on in the mind, which shows that they cannot be identical. The states and processes of your brain are, just like any other bodily states and processes publicly observable, at least in principle. This means that other people may observe them as well as you can. There may well be portions or features of your body that you do not display in public for one reason or another, but others could observe them in principle; in fact some parts of your body can be observed more directly by others than by yourself, such as your eyes, for which you need a mirror, or the back of your head, for which you need two. Your brain and the events in it may never in fact have been observed by anyone at all, but, to whatever extent they could be observed by you, they could be observed equally well (and again, more directly) by others.

In contrast, consider some item in your consciousnessÑsay a pain. Can other people "observe" or experience this pain of yours? They may well

infer that you are in pain from your behaviour, including your linguistic behaviour; if you say "I am in pain," then we would normally infer that you are. Or we might infer that you are in pain from seeing what has just happened to you. A brick has just fallen on your unprotected foot; although you have not flinched or cried out, it is a fair bet that you are in considerable pain. But inferring that you are in pain, however reasonable the inference, is not the same thing as experiencing your pain in the way that you do. Even if neurophysiology develops to the point where the expert can say that such-and-such a brain process is associated with a sharp pain in the left knee, and that since that process is going on right now then you no doubt have such a pain in your left knee, that expert is still only inferring the existence of the pain. To experience your pain in the way that you do, one would have to have itÑand who could possibly have your pain but you? It is said that some twins, or friends or partners, are so close to each other that when one is hurt or in trouble the other (who may be miles away) "feels it"; some men, apparently, feel labour pains when their partners are giving birth. Is such a man, who of course is not himself in labour, feeling someone else's pain, namely his partner's? Not at all. He is feeling only his own pain, even if it is caused in some strange way by hers, or by the physical events in her body, rather than by any obstetric events in his. Thus, it is held, physical things and events are, at least in principle, publicly observable, whereas mental items and events are essentially

private. They are private in the sense that only the "owner" of a mental item can experience it directly, or experience it at all. And this essential difference between the mental and the physical shows that mindbrain materialism cannot be true.

A second important objection, which we shall discuss only briefly, is sometimes called the Infallibility Objection. Suppose you believe that your body is in a certain stateÑfor example, that you have a cut on your hand. This belief of yours, strong though it may be, could conceivably be mistaken: you may be hallucinating, or the "blood" you see may be red dye that someone has spilt on you. There may well be circumstances in which it seems very far-fetched to suppose that such a belief could be mistaken, but it is at least conceivable, and so your belief is said to be fallible. In contrast, take your belief that you are in pain. Could you possibly be wrong about this? What would you think if the physiologist said that he or she had inspected your brain thoroughly, and there is no pain process going on, and so you cannot really be in pain? Would you accept this authority's judgement, and concede that you were wrong? Not at all. You would rightly insist that you are the authority on the question of whether you have a pain. Your belief, that is to say, is infallible. This is not to deny that you could be mistaken as to the source of your pain. You might complain of a pain in your right foot, for example, and have to be told that in fact you do not have a right foot at all, but have lost it, and are suffering

from the "phantom limb" phenomenon in which the nerves that would normally carry the appropriate messages from the foot are being "triggered" somewhere further up the system. It may be too that there are other mental states with respect to which one is not infallible. Being in love is sometimes offered as an example. Could you be mistaken about being in love? How? But at least with respect to some mental states, such as being in pain or experiencing pleasure, it is held that your belief that you are in such a state is infallible. And so, it is concluded, here is a second difference between physical states and at least some mental states, which is another body blow, so to speak, to mind-brain materialism.

Functionalism

The last of the theories we shall introduce, and the newest of them, is one that appeals to many people sympathetic to the idea of machine mentality. Mentality, it is said, consists in certain causal powers possessed by certain things, such as brains, in certain states. A given state of the brain may cause certain things to happen, and another state may cause certain other things to happen; these states are mental states in virtue of the kinds of thing they cause, and they are different mental states (or mental states of different kinds) in virtue of the different kind

of thing they cause. Mental states are defined and classified by what they do, by their functions.

Two important questions arise. First, of what general type, or types, are the functions concerned? Many things have functional states that are not mental states; a lawnmower, for example, has states that together cause the lawn to be mown (engine running, blades attached, mower moving across lawn, etc.) but are not mental states. So what types of function count, for mentality? Second, could anything other than a brain have states that serve the appropriate functions?

On the first question, functionalists differ considerably among themselves. Some specify the functions in terms of outward behaviourÑa given mental state is a state that produces certain behaviour of a kind we normally

associate with that mental state. The state of amusement, for example, is one that produces amusement behaviour such as smiling or laughing. But this runs into the problems we encountered with behaviourism: In some cases there is amusement without amusement behaviour, or vice versa, and the same goes for many other states such as pain fear, anger, belief, and so on. This makes it very difficult to define mental states of particular kinds in terms of what they do; they do not always do it, and other states do it too. Other functionalists seek to solve this problem by saying that the functions in question may include the causing of certain other mental states. So, for example, the mental state of amusement might be said to cause the desire to smile or laugh, as well as actual

smiling or laughing. Such a move may cover the case of the person who is amused but does not show itÑif we can suppose that he or she has at least a desire to smile or laugh, caused by being amusedÑbut there is still the problem of the faker. When you laugh at the rich man's feeble joke you are not only laughing, but doing so because of a desire to laugh. Yet this desire has not been caused by amusement, but by some other mental state such as a desire to impress him. So we have not yet functionally differentiated the mental state of amusement from other mental states such as this.

Setting this problem aside, or supposing that it can be dealt with, an important thing to notice about the appeal to causing other mental states is that we are giving up hope of defining mental states in general. If we have to say that a certain mental state is one that causes certain other mental states, how much have we said about the nature of mental states as such? It seems rather like saying that a horse is something that produces other horses. Does this tell us much about what it is to be a horse? Some functionalists do not mind this too much, for they are more interested in establishing causal links among states of various specific kinds than in defining mentality in nonmentalistic terms. And much of the work they have done is very sophisticated and instructive in extending our understanding of mental operations and relationships. But let us now turn to question (2) above. Whatever the causal functions required for mentality or for any particular instance of it might be, could anything other than a brain furfil them? Could a computer?

Functionalists differ as to whether anything other than a brain will in fact ever do so. But they agree on this: There is no reason in principle why a computer could not. It may require circuitry that copies that of the brain, and such sophisticated electronics may be a long way off, but it cannot be declared impossible. Some functionalists go further and say that even such brainlike physical circuitry may not be required, that even an immaterial substance like Descartes' mind could (conceivably) have such functions. The crucial thing for a functionalist is what a state can do, not what it is a state of. One point that seems to justify the functionalist's openness to the idea of things other than brains having mental states is that the nature of those states, and their relationship to physical states, are at best very mysterious indeed. As we have seen, it is not at all easy to define mental states in physical terms or any other terms, and the "reductionist" attempts by behaviourists and other materialists to show that they are nothing but behavioural dispositions, or brain states, run into serious problems. If the relationship between the material and the mental, the physical and the psychical, is impenetrably mysterious anyway, why should we narrow-mindedly rule out the possibility of such things occurring in computers or even indeed in immaterial spirits, should there be any? Immaterial spirits are somewhat

Mozart resulted from ideas that had never been had (at least in those forms and combinations) before, but those ideas in those forms and combinations were still produced by what had gone before. And any being with enough comprehension of the

relevant circumstances and laws of psychology and/or physiology and so on could have predicted th

artistic productions (together with every other human thought and deed) in detail, and maybe somewhere such a being exists and did predict them. The actual level of our understanding seldom however permits more than rough predictability of some human behaviour, and so we are surprised and impressed when something new and undreamt of is produced.

The view of human nature just sketched is known as determinism. It is a highly controversial view, with some very challenging apparent implications for human dignity and moral responsibility, and around it there still rages a major and manyfaceted philosophical debate. To enter that debate here would unfortunately take us well away from our present concern. We raise the topic not in any hope of resolving it, but to query further the Lady Lovelace challenge to the Turing Test. Any degree of creativ~ty suggestive of mentality may, for all we know, be the result of "programming" in ourselves or in anything else. So the legitimacy of the Test seems to be upheld.

Another challenge, however, comes from John Searle. Remember his parable of the Chinese Room. With the right instructions in the room and enough practice, says Searle, he can fool native Chinese speakers; he can thus pass a test analogous to the Turing Test, but without understanding any Chinese. So, if the computer's performance could also be the result of very good instructions, that is, its program, its ability to pass the Turing Test no more proves its mentality than Searle's ability to fool the Chinese speakers proves that he understands Chinese (1980: 419). In response to this argument it is often said that, although the person in the room does not understand Chinese, the whole system does. The system includes the room, the baskets of symbols, and the ledgers giving the detailed instructions for Searle to follow, as well as Searle himself. And this system as a whole understands, or incorporates an understanding of, Chinese, how else could the instructions be so good as to enable Searle to fool the native speakers? Similarly, the computer system as a whole, including the hardware and the software, is what passes the Turing Test and therefore demonstrates understanding, that is, mentality.

Searle says that this reply is inadequate. Reminding us of his distinction between syntax and semanticsÑbetween following formal procedures, a

he does in the Chinese Room when he follows the instructions, and attaching meanings to those procedures, which he cannot doÑhe says that syntax is not aufficient for semantics, and there is no way that the system ca

get from the former to the latter. And in this he seems to be quite right. But does it follow that the computer could not acquire the semantics in some

other way? Between sessions in the Chinese Room, Searle might enroll for lessons in Chinese and begin to make sense of the messages he is handling in the room. Similarly, could the computer acquire mastery of the semantics in some way other than via its programmed mastery of the syntax?

How could it do this? Let us approach this question by imagining three artefacts, named Tom, Dick, and Harry. Tom, synthesised by a clever biochemist, is molecule-formolecule indistinguishable from a human being. Searle imagines such an artefact, and says that he could presumably think, but is just a "surrogate human being" (1984: 36). Dick is a digital computer, with an impressive program that enables him to pass the Turing Test. This, on Searle's view, would not show that Dick can think. Harry differs from Tom and resembles

Dick in being nonbiological, being made of silicon and plastic and so on. But he differs from Dick in having a brain that duplicates the neural structure of a human being; the neurons and the pathways and the rest are replicated electronically. Suppose that Harry then sails through the Turing Test, answering every question with flair, and lying, where necessary, through his silicon teeth. Could he have acquired semantics, as we do? Searle may be justified in refusing to accept Dick as a thinking computer. But what about Harry?

ROBOT RIGHTS

Many people's first reaction to the prospect of thinking computers will no doubt be one of fear. Will they take over the world? What will they do with us? Many a sci-fi scenario has exploited this fear, describing bizarre fates that may await the human species. Perhaps there will be specific AI projects that should be stopped or disallowed because of such risks; perhaps any attempt to build a thinking machine presents unacceptable risks, whatever the intentions of the builders and programmers might be. Let us suppose that it is too late to stop them, and we have before us Harry. Suppose too that there seems no reason to fear him (it?), or at any rate no reason to fear it any more than the average stranger you meet in the street. How should we treat it? The stranger in the street has various rights, many of which have implications for your conduct; you ought not, for example, rob or assault or kill him or her, and maybe you ought not to do certain things that he or she may find threatening or offensive. What about Harry? Does it have rights too? Rights to what?

Some of the particular rights possessed by a human seem hardly applicable to a computer. Not being alive it cannot be killed, and so it has no right not to be killed; and while it could be physically assaulted, it could not be sexually assaulted. But what if Harry has sexual organs as well as a "neural" structure that mimics the human one in such a way as to allow sexual trauma? As far as killing is concerned, it might be suggested that switching off the power supply to a thinking and feeling computer would

stop it from thinking and feeling, and thus be tantamount to killing it. Or, since the power supply may be restored and with it the mental activity, would it be more like rendering a person unconscious? This of course would still be regarded as contravening the person's rights, unless of course it is a case of medical anaesthesia to which the patient or his or her representative has consented; and it may be not far removed from killing if the victim is rendered unconscious for a very long time. But even if some human rights are inapplicable in any form to a computer or robot, it does not follow that it has no rights at all. The right to secondary-level education is not applicable to people who are severely mentally retarded, but they have plenty of other rights.

The retarded person is at least human. Can something that is not human at all be said to have rights? People do talk of animals having rights, but can this really make sense? Animals cannot claim these supposed rig~s, nor are tney seen as nav1ng moral responsibilities, and surely to have rights you must be able to claim them and you must have responsibilities too. Several replies to this may be offered. First, seriously retarded humans cannot claim their rights and cannot be seen as having moral responsibilities, but surely they have the right not to be killed and not to be exploited in various ways. Second, whatever we think about animals who says that a suitably sophisticated computer like Harry could not claim rights or have responsibilities? Third, even if something cannot be said to have rights it may still be wrong to treat it in certain ways: It may be thought wrong, for example, to cut down a magnificent tree or demolish

a beautiful building, and not only because people have a right to enjoy it. It has value, one might say, even if no rights. And perhaps this could apply to Harry, though the second reply seems more relevant.

At this point a rather different move might be made against Harry. Whatever its qualities and abilities might be, it is a human creation. And the creator of a thing has the right to do with it whatever he or she likes if you build a house (out of your own materials) you are entitled to demolish it, unless you have contracted to sell it, and the same goes for anything else you make. So, while nobody else may have the right to destroy or "mistreat" the computer you have built, you may do what you like with it. In response to this we might say that parents, for example, who create children out of their own materials (eggs and sperm), do not thereby acquire the right to do what they like with them. A few people do think they have a right to kill their unwanted infants, and a lot of people used to think that it was quite in order to subject their young children to labour in chimneys or mines. Some people still treat their children as items of property, according them little respect as persons; but even if many of us now are sometimes guilty on this count, we recognise important 1imits on parents' rights over their children. Why should the same not apply to computer engineers, rights over their creations? It is true that

parents do not design their children in the way that computer builders do; the procreative act may not in fact have been intended to be procreative at all. But this difference may be only one of degree. Some parents do try to design their children in certain ways, by choosing an appropriate coparent, adopting certain dietary and environmental habits during pregnancy, and so on; do they thereby acquire greater rights over the life of the child? If by genetic and / or environmental engineering they could have a child designed precisely to order, would they acquire greater rights still? Surely not. The child's potentialities and needs are the basis of important rights that severely restrict the rights of even its parents, whatever the extent to which the parents designed and made the child to be as it is.

There seems, then, no conclusive reason to deny appropriate rights, and the appropriate treatment, to a computer like Harry. We leave it to you to imagine what life with Harry might be like. We also leave you with this question. We may well have difficulty in believing that Harry really does have thoughts and feelings and therefore rights. He (no, it!) is only a machine, after all; it is not made of flesh and blood, but of silicon and plastic and so on. Can it really think and feel, however plausibly it claims to be able to? Many of us may well have a strong and lingering doubt. But given what is at stakeÑlife, or something like it, liberty in an appropriate form, and the pursuit of electronic happinessÑshould we not give Harry the benefit of the doubt?

SUMMARY

In this chapter the question was examined of whether computers, or any machines, could be thinking and feeling things and therefore have moral rights. The discussion began with a consideration of the mind-body problem. There was an attempt to establish that no sound objections exist in principle to machines having minds. Turing's test for intelligence was then examined, and it was suggested that if any machine, or robot, could pass it, we should give it the benefit of the doubt and assume that it does have moral rights.

REFERENCES

Descartes, Rene. 1642. Meditations on First Philosophy; page citations to

edition of Anscombe and Geach, 1971.

Evans, Christopher. 1992. Can a machine think? In Philosophy and Contemporary Issues. 6th ed. Ed. John R. Burr and Milton Goldinger. New York: Macmillan, pp. 423-35.

Hobbes, Thomas. 1651. Leviathan; page citations to the edition of Macpherson, 1975.

Searle, John R 1980. Minds, brains, and programs. The Behavioral and Brain Sciences 3: 41724.

Searle, John R. 1984. Minds, Brains, and Science. London: British Broadcasting Corporation.

Turing, A. M. 1950. Computing machinery and intelligence. Mind 59: 433~0.

FURTHER READING

Clarke, Roger. 1993. Asimov's laws of robotics: Implications for information technology, Part 1. Computer 26 (December): 53-61.

Clarke, Roger. 1994. Asimov's laws of robotics: Implications for information technology, Part 2. Computer 27 (January): 57-66.

Dreyfus, Hubert L. What Computers Still Can't Do: A Critique of Artificial Intelligence. Cambridge, MA: MIT Press.

Epstein, Robert. 1992. The quest for the thinking computer. AI Magazine 13 (Summer): 81-95.

Feigenbaum, Edward A. 1996. Turing award lecture: How the "What" became the "How." Communications of the ACM 39 (May): 97-104.

Haugeland, John. 1987. Artificial Intelligence: The Very Idea. Cambridge, MA: MIT Press.

Johnson-Laird, P. N. 1988. The Computer and the Mind: An Introduction to Cognitive Science. London: Fontana Press.

Matthews, Robert. 1994. Computers at the dawn of creativity. New Scientist 144 (10 December): 30-34.

Reddy Raj. 1996. Turing award lecture: To dream the impossible dream. Communications of the ACM 39 (May): 105-13.

Glossary

AbsolutismÑthe view that actions of certain kinds are always right or always wrong.

Artificial intelligenceÑhumanlike intelligence in machines or the discipline of attempting to develop it.

BehaviourismÑthe view that mental states and events may be understood or defined purely in terms of outward physical behaviour.

BugÑan error in a computer program.

ConsequentialismÑthe normative view that the rightness or wrongness of actions is determined solely by their (likely) consequences. Utilitarianism is a type of consequentialism.

DeterminismÑthe view that everything that happens (including all human thought and behaviour) is caused by what has gone before and is at least in principle predictable.

Diversity problemÑin ethics, the problem (for objectivists) of explaining the diversity in moral beliefs.

Divine command theoryÑthe objectivist theory that the right action is that which is commanded or approved by God.

DualismÑthe view that the mind is an immaterial thing, distinct from the body even if causally related to it.

FunctionalismÑthe view that mental states are to be understood in terms of their causal functions and that in principle things such as computers could have such states.

Greatest happiness principleÑsee Utilitarianism.

HedonismÑthe normative ethical view that pleasure is the only thing good for its own sake.

Image manipulationÑthe alteration of images stored electronically, usually on a computer disk.

InternetÑthe network linking computers worldwide.

IntuitionismÑa type of objectivism stating that we know moral truths by a special intellectual faculty, intuition.

LambdaMOOÑthe most famous MOO, and primarily the context for a game in which players take on personalities and interact in various ways with one another.

Meta-ethicsÑthe study of the status or meaning of moral beliefs.

Mind-brain materialismÑthe view that mental states and events are identical to brain states and events.

MOOÑa "place" or environment on the Internet where people can interact online. Currently MOOs are mainly used for game playing.

NaturalismÑthe meta-ethical view that morality may be defined in terms of ordinarily observable things and qualities, as opposed to supernatural ones.

Normative ethicsÑthe study of particular moral questions, and of general principles and stances on how we ought to act.

ObjectivismÑthe meta-ethical theory that there are objective moral truths independent of societies and individuals.

Other-regarding actionsÑactions that may (sign)ficantly) affect other people, as opposed to self-regarding actions.

PaternalismÑrestricting the liberty of someone for his or her own good.

PluralismÑthe normative ethical view that our duties cannot be reduced to one fundamental principle.

PrologÑa computer programming language based Qn logic (Programming in Logic).

RelativismÑthe meta-ethical position that there are no objective moral truths. All moral beliefs are relative to, for example, a culture (cultural relativism) or the individual (subjectivism).

Self-regarding actionsÑsee Other-regarding actions.

SubjectivismÑthe meta-ethical relativist position that moral beliefs are relative to the individual; for example, the view that moral judgements express the speaker's desires and aversions.

Turing TestÑa suggested test for art)ficial intelligence in terms of the machine's ability to converse.

UtilitarianismÑthe normative ethical view, closely identified with Jeremy Bentham and J. S. Mill, that the right action is that which promotes the greatest happiness or pleasure for the greatest number of people affected.

Verification problemÑin ethics, the problem of establishing the truth or falsity of supposedly objective moral judgements.