

Ethics, Speculation, and Values

Rebecca Roache

Received: 29 October 2008 / Accepted: 29 October 2008 / Published online: 18 November 2008
© Springer Science + Business Media B.V. 2008

Abstract Some writers claim that ethicists involved in assessing future technologies like nanotechnology and human enhancement devote too much time to debating issues that may or may not arise, at the expense of addressing more urgent, current issues. This practice has been claimed to squander the scarce and valuable resource of ethical concern. I assess this view, and consider some alternatives to ‘speculative ethics’ that have been put forward. I argue that attempting to restrict ethical debate so as to avoid considering unacceptably speculative scenarios would not only leave scientific progress devoid of ethical guidance, but would also rule out some of our most important ethical projects. I conclude that the issue of speculation is a red herring: what is most important is not that ethicists concentrate on current issues or those that are most likely to arise; but that ethicists, scientists, and others focus on maximising what is most valuable.

Keywords Cost-benefit analysis · Enhancement · Ethics · Nanotechnology · Precautionary principle · Science · Speculation · Technology · Values

Discussion of human enhancement technology and nanotechnology is very often future-oriented. Whilst there currently exist drugs, tools, and techniques to raise our capacities above the level of what we might loosely call normal functioning, and whilst scientists are currently able to perform certain operations at the nanoscale, it is by considering how such technologies might develop in the future that we encounter more exciting prospects for transforming ourselves and the world, and more serious ethical issues. It should not be surprising how easily such discussions move from considering what is currently possible to considering what might become possible: philosophers, in particular, regularly take what is currently technically possible one step further in order to clarify the philosophical issues at stake. For example, Derek Parfit moves from the observation that certain brain surgery patients have ended up with a ‘divided mind’ to considering the possibility of dividing the brain and transplanting the two halves into separate bodies, in order to examine common intuitions about personal identity and self-concern [1]. And, Daniel Dennett moves from knowledge about the ways in which sensory information and motor signals are communicated to and from the brain to considering the possibility of communicating such information remotely, in order to question common intuitions about the nature of the self [2].

R. Roache (✉)
Future of Humanity Institute, University of Oxford,
Suite 8, Littlegate House, 16-17 St Ebbe’s Street,
Oxford OX1 1PT, UK
e-mail: rebecca.roache@philosophy.ox.ac.uk

It is not only philosophers, however, who consider such hypothetical possibilities. Public debate in the media is rife with discussion about what might result from embarking on a slippery slope by endorsing a particular piece of legislation. For example, the following concerns are familiar from recent public discussion: if the laws on using cannabis are relaxed, then more people could be drawn in to using harder drugs; if embryos are created for research purposes, then we will be encouraged to view humans as disposable; and if we increase the amount of time that the police are allowed to detain suspects without charge, then further incursions on our freedom can be expected in the future. These hypothetical cases differ in an important way from those imagined by Parfit and Dennett. Parfit and Dennett do not claim that the scenarios they imagine are ever likely to happen; the value of considering them is due to the fact that they help us examine our intuitions.¹ On the other hand, any value of considering the scenarios claimed to lie at the bottom of a slippery slope on which we are about to embark derives from the belief that they might actually happen. Such scenarios, we might say, are not merely hypothetical—they are speculative. We take some such speculative scenarios seriously, in that we deem it worth preparing for them, either by taking steps to avert them or by equipping ourselves to deal with them. But, can we be relied upon to judge correctly which speculative scenarios we should take seriously? And, how can we best strike a balance between preparing for merely possible future eventualities and addressing those issues that are currently most pressing?

Getting this balance right is not easy. It is particularly difficult when considering how to prepare for future technologies like nanotechnology and human enhancement, which many claim could have

¹ Some have disputed the value of considering such hypothetical scenarios. Ludwig Wittgenstein believed that exotic thought experiments can be misleading and meaningless: ‘It is as if our concepts involve a scaffolding of facts ... If you imagine certain facts otherwise ... then you can no longer imagine the application of certain concepts’ ([3]: proposition 350). And Alasdair Urquhart has accused discussions in the philosophy of mind and cognitive science of a ‘fascination with far-fetched thought experiments’ ([4]: 27, cited at [5]: 34f), which illicitly derive their power from ignorance of the underlying science. For an illuminating discussion of the value of thought experiments, see [6].

an unprecedented effect on our lives, and on what it means to be human. That their effects could be unprecedented means both that it is important to consider antecedently what ethical constraints we might want to place on the development and use of these technologies, and that past experience is likely to be an inadequate guide to what we can expect from them in the future. As a result, discussions about the ethics of nanotechnology and human enhancement are replete with speculation. This is frowned upon by many writers, who complain that highly speculative scenarios receive too much attention in such discussions, and suggest measures to divert ethical attention away from speculation and towards currently pressing issues. The extent to which speculation is acceptable in ethics has not been addressed head-on, however, and this topic will be the focus of this paper. Taking as a case-study ethical debate about nanotechnology and human enhancement, I will argue that not only is engaging with speculative scenarios often appropriate in ethics, it is very often necessary. Further, attempts to restrict the extent to which ethicists² deal with speculative scenarios risk ruling out many of what seem intuitively to be correct and sensible decisions and projects. As such, I aim to explore issues that might inform debate about enhancement, and which might also be of broader theoretical interest.

The Case Against Speculative Ethics

Several writers have highlighted problems that arise from structuring ethical debate around speculative, sometimes incredible scenarios. For example, Neil Manson warns that giving too much weight to highly speculative possibilities in decision-making has led to certain groups adopting an incoherent version of the so-called ‘precautionary principle’ [7]. Adam Keiper accuses participants in the debate about the ethics of nanotechnology of addressing certain issues ‘prematurely’ ([8]: 55): ‘[i]n much of the burgeoning nanotechnology literature, there is a sloppy and lazy tendency to slip from today’s cutting-edge science to the most far-

² I use this term broadly, to include not only academic or philosophically-trained ethicists, but people in general whose job it is to decide how best to prepare for future scientific developments. This is in line with the concerns of the authors whose views I consider here.

out imaginings of futurists, as though the former were old news and the latter were ineluctable' ([8]: 64). And, so serious and widespread does one writer view the tendency to devote ethical attention to possibly spurious future scenarios that he has designated it a 'syndrome'. In an earlier issue of *Nanoethics*, Alfred Nordmann claims that 'speculative ethics'—as he terms the disproportionate focus of ethical attention on highly speculative scenarios³—suffers from

a radical foreshortening of the conditional, that is, ... what one might call the 'if and then' syndrome. An if-and-then statement opens by suggesting a possible technological development and continues with a consequence that demands immediate attention. What looks like an improbable, merely possible future in the first half of the sentence, appears in the second half as something inevitable. And, as the hypothetical gets displaced by a supposed actual, an imagined future overwhelms the present. ([5]: 32)

One example of such objectionable foreshortening occurs in discussions of radical life-extension. Thanks to advocates like Aubrey de Grey,

[t]he true and perfectly legitimate conditional 'if we ever were in the position to conquer the natural ageing process and become immortal, then we would face the question whether withholding immortality is tantamount to murder' becomes foreshortened to 'if you call into question that biomedical research can bring about immortality within some relevant period of time, you are complicit with murder'⁴—no matter how remote the possibility that such

research might succeed, we are morally obliged to support it. ([5]: 33)

The main problem with the 'if and then' syndrome, according to Nordmann, is that it leads to the scarce resource of ethical concern being 'squandered on incredible futures' and thereby being 'distract[ed] from ongoing developments that demand our attention' ([5]: 34).⁵ Others writers have expressed related concerns. For example, Keiper complains that the belief that 'ethics has to "keep up" or "catch up" or "evolve" with advances in technology—is a prescription for a shallow and reactive ethics, one that ignores the questions that matter most' ([8]: 67). Cass Sunstein, responding to the claim that it is appropriate for governments to regulate to avoid scenarios that may or may not arise rather than to proceed according to cost-benefit analysis, argues that failing to adhere to cost-benefit analysis results in 'a high level of inefficiency', in which precious resources often fail to be directed towards the most useful ends [10]. And, John Seely Brown and Paul Duguid comment that Bill Joy's [11] speculative fears about where new technologies might lead us are 'a form of tunnel vision, excluding broader social responses. Rather than arming society for the struggle, the debate may not only be alarming society, but unintentionally disarming it with a pervasive sense of inescapable gloom' [12].

So, when is a prediction too speculative for ethicists to waste their time considering? We have seen that considering speculative future possibilities comes at the expense of addressing pressing current issues. However, we cannot do away completely with considering speculative future possibilities, since doing so would leave us unable to address pressing current issues. Consider that any decision about which course of action to take will involve attending to possible future scenarios. For example, my rela-

³ Others have used the term 'speculative ethics' to refer to other things. For example, John Maynard Keynes used it to refer to the investigation of what is good in and of itself (Keynes JM (1905) *Miscellanea Ethica*. Unpublished manuscript deposited in King's College Library, University of Cambridge), and Albert William Levi, based on Henri Bergson's distinction between knowing a thing by moving around an object and knowing a thing by entering into it, used it to refer to a method of ethics that proceeds by 'moving around' moral phenomena rather than attempting to penetrate them [9].

⁴ Nordmann does not attribute this particular claim to de Grey, or to anyone else. He appears to use it only to illustrate his claim that ethical deliberation can escalate into a demand for action.

⁵ He lists additional reasons why the 'if and then' syndrome is worrying: it casts purely speculative scenarios as more probable than they really are, or as more probable than we have reason to believe they are; and it prompts questions about ourselves and our future that are 'unintelligible'. I shall ignore these other concerns, however. By themselves, they do not constitute a compelling case against speculation in ethics, and—for reasons I do not have space to outline here—the latter problem arguably does not exist at all.

tively trivial decision about whether to make myself a cup of tea in the next 5 min involves cogitations that, were they fully conscious, might be expressed along the lines of, ‘If I don’t make a cup of tea in the next 5 min, then I will probably grow parched and irritable’, and, ‘If I make a cup of tea in the next 5 min, then I may lose the thread of the paragraph I’m currently writing and will need to start again’. For all I know, the future scenarios to which I attend when deciding whether to make a cup of tea may be no more or less probable than those to which we attend when predicting the long-term prospects for human enhancement, and yet we all attend to such scenarios many times each day, without believing that, in general, doing so distracts us from more important activities. Further, since any decision about which course of action to take involves attending to possible future scenarios, and addressing current issues involves—among other things—making decisions about which course of action to take, attending to possible future scenarios can hardly be thought to take place at the expense of addressing current issues. Rather, attending to possible future scenarios is an important part of addressing current issues.

Why, then, do some writers believe that considering certain future scenarios is a waste of time for ethicists? A promising answer is that ethicists may be liable to overestimate the probability of certain such scenarios. This worry lies behind Keiper’s remark that ethicists often debate certain issues in nanotechnology ‘prematurely’, and Nordmann’s complaint that, in a ‘foreshortened’ conditional, ‘[w]hat looks like an improbable, merely possible future in the first half of the sentence, appears in the second half as something inevitable’. Since, other things being equal, we prioritise those issues that are ‘inevitable’ over those issues that are less likely to arise, overestimating the probability of an event will lead us to overestimate the urgency of ethical issues relating to it. And, overestimating the urgency of certain ethical issues will lead to an inefficient application of ethical expertise.

What explains ethicists’ alleged tendency to overestimate the accuracy of certain predictions? The answer may be that those who make such predictions often play down the extent to which they are highly speculative and incredible. Brown and Duguid observe that, despite the fact that accurately predicting the future direction of technology is a tricky business,

requiring one to take into account not only technological capabilities but also a variety of social factors,

What we tend to get [instead] is the simpler kind of extrapolation, where the path to the future is mapped out along vectors read off from technology in isolation. Following these vectors, it’s easy to count in the order of ‘1, 2, 3,... one million,’ as if intervening steps could be taken for granted. So unsurprisingly a post-bomb book from 1945, written when no one had even developed a nuclear car engine, notes that[, ‘] Production of the atomic-energy type of motor car will not entail very difficult problems for the automobile manufacturer... it will be necessary to replace the 30,000,000 now registered in a few years.[’] Elsewhere, nuclear energy experts were predicting as late as 1974 that, spurred by the oil crisis, some 4,000 U.S. nuclear power plants would be on line by the end of the century. (The current figure [in 2001 is] around 100, with no new ones in production.) And [in 1940, the Cal. Tech physicist] Langer strides from the bomb to the [uranium]-235-powered house with similar ease, claiming blithely, ‘None of the things mentioned has yet been worked out, but the difficulties are difficulties of detail.’ [12]

It is not hard to see that, if scientists overestimate the likelihood of their technological predictions, then ethicists—who are much less well-qualified than scientists to evaluate the likelihood of such predictions—may end up overestimating the urgency of addressing ethical issues surrounding the predictions in question.

Another possibility is that some writers may be influenced by the affect heuristic. Using this heuristic, the perceived desirability or undesirability of a possible scenario shapes people’s judgments about how likely it is that the scenario will occur [13]. As a result, lay judgments of an outcome as desirable or the reverse are correlated with assessments of, respectively, lower and higher risks of the outcome’s occurring⁶. This may lead some ethicists to overesti-

⁶ In fact, the picture is more complex than this. Judgments about the likelihood of a given scenario’s occurring depend on a variety of factors, of which its desirability is only one. Other relevant factors include whether the activity that might lead to the scenario is voluntary, whether it is perceived as controllable, and whether it is familiar. See [14].

mate the likelihood, and therefore the urgency, of certain undesirable outcomes.

Rooting out Speculation

If debating highly speculative scenarios is a waste of ethical concern, how can we best select which ethical issues to address when considering novel technologies? Nordmann outlines two strategies for doing this, which he suggests as alternatives to the approach characterised by the objectionable ‘if and then’ syndrome. According to the first, one ‘confine[s] oneself to just those ethical issues that are present already in the process of emergence and thus prior to the appearance of any technical novelty’ ([5]: 34). This involves addressing issues like technical possibility, social desirability, and ‘the visions that are promoted by funding calls and applications’, as well as more general issues like honesty and distributive justice. According to the second strategy, one ‘exercises the virtue of patience and enters the debate only as particular issues present themselves’; for example, ‘[a]fter some time of waiting for a genuine nanotechnological problem, the scientific and technical questions surrounding nanoparticle toxicity have finally given rise to a new discussion of risk governance, of epistemic and objective uncertainty’ ([5]: 34). Nordmann goes on to comment that

Rather than adopt a believing attitude towards the future, an ethics beholden to present capabilities, needs, problems, and proposed solutions will begin with vision assessment. Envisioned technologies are viewed as incursions on the present and will be judged as to their likelihood and merit: How credible are these claims, and do these technologies solve acknowledged problems? More generally: What do these visions tell us about the present, what is their implicit criticism of it, how and why do they require us to change? ([5]: 41)

Michael Crow and Daniel Sarewitz make a similar suggestion. Having remarked that, whilst it is reasonable for ethicists to consider predictions arising from ‘simple extrapolations of current trends in technological innovation and societal transformation[, m]ore adventurous speculation is ... best confined to science

fiction novels’ ([15]: 97), they claim that, in order to be prepared for whatever social effects nanotechnology might have, we need to develop ‘knowledge and tools for more effectively connecting R&D [research and development] inputs with desired social outcomes’ ([15]: 99). This might involve ‘[a]nalysis of past and current societal responses to transforming technologies’ ([15]: 99), ‘[c]omprehensive, real time assessment and monitoring of the nanoscience and nanotechnology enterprise’ ([15]: 100), ‘[a] science communication initiative, to foster dialogue among scientists, technologists, policy makers, the media, and the public’ ([15]: 100), and ‘[a] constructive technology assessment process, with participants drawn from representatives of the R&D effort, the policy world, and the public’ ([15]: 100).

John Weckert and James Moor, too, counsel against considering highly speculative scenarios. They suggest that, when using the precautionary principle to decide how to prepare for developments in nanotechnology, ethicists should focus only on ‘credible’ threats—that is, those such that ‘an hypothesis that that threat is caused by a particular action is a reasonable hypothesis, even if there is little actual evidence to support the causal link’ ([15]: 143). This has the result that, whilst some familiar worries about nanotechnology are deemed worthy of our consideration, others turn out to be a waste of time:

The hypothesis that nanoparticles cause harm to humans is reasonable given what is known about asbestos and deserves further testing[.] It is plausible to believe that they might be harmful even though there is not enough evidence to even say that this is probable. It is less clear that gray goo presents a credible threat [since] there are serious doubts about whether self-replication of the type required is possible. If this is so, then an hypothesis such as ‘the development of self-replicating robots will lead to the gray goo problem,’ while perhaps true, is practically pointless given that the development of these robots is such a remote possibility. ([16]: 143–144).

All of these suggested strategies emphasise that ethicists should focus on responding and reacting to scientific innovation, and on debating issues arising from what there is scientific reason to believe probable, rather than on ethically evaluating highly speculative scenarios.

It is important to note that, whilst Nordmann, Crow and Sarewitz, and Weckert and Moor advise against wild speculation by ethicists, they do not criticise scientists for engaging in it during the process of innovation.⁷ A moment's consideration makes it clear that they are wise not to do so. The most exciting future technologies are novel and often achieve what, previously, was barely dreamed possible: consider space travel, *in vitro* fertilisation, or even the explosion in computer processing power. Such achievements can be realised only by questioning current assumptions about what is possible, and perhaps by adopting the sentiment once expressed by George Bernard Shaw: 'You see things: And you say, "Why?" But I dream things that never were: And I say, "Why not?"' ([17]: part I, act I). What is *de rigueur* for innovation in science, however, is—according to some—unacceptable for ethicists. But, is it sensible, or even possible, to confine ethical consideration of technology to current and emerging technology?

I will suggest several reasons why it is either impossible or imprudent to restrict ethical debate in this way. First, it risks squandering scientific effort and making it more difficult to contain unethical scientific projects. Second, highly speculative and sometimes improbable future scenarios are sometimes instrumental in motivating important ethical projects. Third, the claim that it is always inappropriate to take action to prepare for such highly speculative and sometimes improbable future scenarios is too strong, and contradicts some important ethical intuitions. Finally, the emphasis that some writers place on addressing currently pressing issues may distract us from what is really important.

⁷ Occasionally Nordmann seems to do this. For example, he writes, 'In order to resist foreshortening, considerable work is required to hold the scientific community to its own standards of honesty and clarity. Whose responsibility is it ... to remind scientists ... of the categorical difference between a therapeutic brain-machine interface and the vision of a thought-controlled mind-machine interface?' ([5]: 43). However, we can interpret him to mean not that scientists should avoid speculation—that is, imagining what might be possible given the right technology—but that it is important that scientists should be clear about what is involved in achieving those scenarios that they envision.

The Case for Speculation

We have seen that some writers are concerned that—to use Nordmann's words—'ethical concern is a scarce resource and must not be squandered on incredible futures, especially when they distract from ongoing developments that demand our attention' ([5]: 34). By the same token, we might plausibly claim that scientific expertise is a scarce resource and must not be squandered on developing technological visions that are premised on implausible ethical assumptions, especially when doing so distracts us from using science to solve genuine ethical problems. Consider that many technological developments arise as a response to ethical concerns. For example, carbon capture and storage technology is a response to ethical concerns about global warming, and preimplantation genetic diagnosis technology is a response to ethical concerns about creating children with inherited diseases. Identifying an ethical issue, judging it sufficiently serious to warrant a technological response, and deciding what form that technological response should take is a process that requires both ethical and scientific input. Further, the ethical input is required from the very beginning.

The views that we considered in the previous section, however, would have ethicists involved at a much later stage: they would not be involved in identifying the problem and formulating the solution, but merely in evaluating the solutions envisaged by scientists. This criticism applies even to those who—like Crow and Sarewitz—advocate dialogue between scientists and ethicists. Provided that scientists are allowed free reign to speculate (which, as we have seen, they must be) whilst ethicists must confine their attention to current or emerging technology (which, as we have seen, Crow and Sarewitz and others advocate), it turns out that whereas scientists are bound only by their imaginations in choosing the ends at which they aim, ethicists are restricted to evaluating whichever of the scientists' projects turn out to be promising enough to warrant their consideration. This raises the risk of scientists squandering their time developing misguided solutions to ethical problems that they are ill-equipped to interpret and assess. For example, if ethical concern about global warming is the chief motivation for creating carbon capture and storage technology, and if it turns out that such ethical concern is misconceived, then there remains little

motivation for developing and using such technology. However, if it is deemed a waste of ethicists' time for them to evaluate such technology before it is so advanced that its prospects of eventual success are sufficiently high, much scientific effort will have been squandered in developing the technology before ethicists condemn it. Because of this, ethical evaluation of proposed scientific projects at a very early stage—even before their viability has been assessed—is necessary to help ensure that scientific effort is directed at the best ends.

There is another, related reason for resisting the proposal to restrict ethical discussion about technology to current or emerging technology. Plausibly, ethically misguided or unethical technological projects are more easily avoided if such projects are rooted out at a very early stage, before scientists' time, money, and careers have been invested in them. If ethicists wait until unethical technology is current or emerging before criticising it, it will be much more difficult to prevent its development, since the scientists involved will have an interest in ensuring its completion in order to bring to fruition efforts already expended. Much of the already scarce ethical concern would then need to be directed at battling with scientists over whether or not certain projects should go ahead. Attempting to ensure that ethical concern is not squandered by restricting the domain of acceptable ethical consideration, then, could result in an inefficient use of both ethical concern and scientific expertise.

Second, without giving serious ethical attention to 'adventurous speculation' ([15]: 97) about 'improbable, merely possible futures' ([5]: 32), many of our most important ethical projects would never even have been entertained. The quest to develop a cure for cancer, for example, was borne despite lack of knowledge about whether such a cure could ever exist. Rather, it was borne of the speculative belief that a cure for cancer *might* be possible, combined with the ethical evaluation that a cure for cancer would be a great good. Whereas the writers cited in the previous section urge that ethical debate about technology should respond to what is currently technically possible, here we find the process working in the opposite direction: the attempt to render a solution technically possible is motivated by the ethical judgment that such a solution would be a great good. The views cited leave no room for ethical

judgments about what speculative goods might be desirable to determine the direction of technology. And, if the direction of technology is to be determined by non-ethical considerations, we may well worry about where it may take us.

Third, the claim that ethical concern is squandered when applied to scenarios that are not 'credible' ([5]: 34, [16]: 143) is highly controversial, and may well be misconceived. That humans are often not very good at judging the probability of a given scenario is well-documented (see, for example, [18]), meaning that to dismiss certain possible futures due to their incredibility is a risky strategy: some of them could turn out to be much more likely than we currently think. We might circumvent this epistemological issue by reformulating the claim so as to urge ethicists not to squander their concern on what are *known* to be *highly unlikely* futures. As such, however, this claim is controversial. Consider that a popular way of deciding on a course of action is to conduct a cost-benefit analysis. We evaluate the expected costs and benefits of each possible course of action, weight them by their probabilities, and pursue the course of action that promises the best net benefits. According to this method, that a possible future event or scenario is highly unlikely is not sufficient reason to exclude it from our consideration, since our decision must take into account not only an event's likelihood, but also its value—that is, the costs and benefits it brings. Some highly unlikely future events nevertheless have consequences that are so bad that the course of action with the best net benefits involves taking measures to prepare for such events. For example, the extremely slim chance that the Large Hadron Collider—the world's biggest particle collider, currently under construction—will result in the extremely bad consequence of destroying the earth has led to a lawsuit to halt its creation [19]. Or, to take a less exotic example, imagine that a child has become lost in a large network of underground caves, and that—since the child is not an experienced caver—the chances of him having a fatal accident are high. A rescue team of experienced cavers is available to search for the child. It may be highly unlikely that the rescue team will be able to reach the child alive, but even so, the slim chance of saving the child's life is nevertheless a large enough good to outweigh the costs (in terms of time, equipment, the risk of injury to the members of the rescue team) of conducting the rescue operation.

Indeed, we may well react with horror to the suggestion that the rescuers should not even attempt to look for the child because the chances of finding him are too low. Expending effort on considering incredible or highly unlikely futures, then, does not necessarily squander ethical concern: whether such efforts are worthwhile depends not only on the likelihood of the futures, but also on the nature of the costs and benefits that arise if those futures materialise.⁸

We saw earlier that Nordmann criticises de Grey for his insistence that we address ethical issues relating to the speculative possibility that scientists may 1 day be able to halt the ageing process. Considering how we might apply cost-benefit analysis to deciding whether or not to address these issues gives us an interesting insight into the claim that ethicists squander their time when debating highly speculative scenarios. Is it possible, at least in some cases, that disagreements about whether or not certain possible future scenarios warrant ethical attention stem not from divergent views about how probable *in general* future scenarios must be in order to warrant ethical attention, but from divergent views about the value of those scenarios? For example, Nordmann objects to de Grey's suggestion that we are morally obliged to support attempts to conquer the ageing process even though there are currently no compelling reasons to believe that such attempts will succeed. It seems less likely that de Grey believes that we should *in general* pursue ends that we have inadequate reasons to believe obtainable, than that he places a greater value on conquering ageing than Nordmann does. As a result, according to de Grey's calculations, cost-benefit analysis dictates that we pursue the goal of defeating ageing

despite the absence of compelling reasons to believe that we will succeed, and so addressing the ethical implications of this is worthwhile. According to Nordmann's calculations, cost-benefit analysis dictates that we do not pursue this goal, and that addressing its ethical implications is not worthwhile. Another example from Nordmann: he criticises Nick Bostrom for arguing that, since there is a non-negligible possibility that intelligent machines will be built within 50 years, it is 'important to consider' this possibility ([20]: 41, cited at [5]: 39). Nordmann accuses Bostrom of attempting to reverse the usual burden of proof by adhering to the implausible principle that '[i]f we can't be sure that something is impossible, this is sufficient reason to take its possibility seriously' ([5]: 39). Few would seriously endorse this principle; however, if we believe that a certain potential danger is grave enough, we may be moved to prepare for it. We might, then, take Nordmann's disagreement with Bostrom here to stem from a disagreement about the gravity of the potential danger of intelligent machines.

Whilst the critics of speculative ethics that we have considered explicitly state that it is *speculation* in ethics to which they object, taking their objection to relate also to the value of the outcomes debated would enable them to avoid some controversial implications of their views. They would not, for example, be committed to denying ethicists a role in shaping the earliest stages of future technology; rather, speculative future technologies may constitute worthy topics of ethical debate provided that they promise to deliver sufficiently valuable goods if they do turn out to be tenable.

Cost-Benefit Analysis on Trial

Pursuing purely speculative goals, and the method of cost-benefit analysis, do not escape scot-free, however. Rightly or wrongly, for pragmatic reasons there are limits to the efforts we are willing to expend pursuing extremely small chances of securing extremely great goods. There is presumably some tiny chance that allowing hundreds of cats to roam around a scientific laboratory will lead to their accidentally and dispassionately creating a potion that turns out to cure all of the world's diseases, and yet even the possibility of this enormous good is

⁸ It may be objected that Weckert and Moor are concerned with the precautionary principle, which is often said to be an alternative to cost-benefit analysis, and as such the observation that their view is at odds with cost-benefit analysis can hardly be deemed an objection. However, it is worth noting that, despite the fact that those who advocate using the precautionary principle rather than cost-benefit analysis do so because the former is more risk-averse than the latter, the version of the precautionary principle that Weckert and Moor favour—which considers only 'credible' threats—fails to consider some risks that even cost-benefit analysis would recognise. If the costs of a particular outcome are sufficiently severe, cost-benefit analysis advocates preparing for that outcome even if it is highly unlikely.

(I assume) unlikely to persuade any funding council to finance such a project. In general, we shun projects that offer miniscule chances of securing extremely great goods in favour of investing in projects that offer more promising chances of securing more modest goods. As we have seen, however, we do sometimes deem it worthwhile to invest in projects whose chances of success are slim, and so the line between worthwhile, promising projects and unacceptably speculative projects is difficult to draw. On the one hand, we make use of cost-benefit analysis to decide which projects to pursue; and on the other, there seems to be a limit to the amount of improbability we are willing to tolerate, regardless of the value of the good we stand to secure if the dice roll in our favour.

Since it is unclear where the division lies between worthwhile projects and unacceptably speculative projects, there exists the danger of a slippery slope, with increasing resources gradually being directed at increasingly unacceptable projects. From this point of view, it is highly relevant to worry that the attention paid to some speculative projects may be undeserved, and may divert resources from other, more promising projects. However, the measures that some advocate introducing to prevent this happening take us too far in the opposite direction. Bostrom observes that

many ... argue that the S&T [science and technology] enterprise needs much more guidance from society in order to ensure that scientific and technological research is really directed to achieve socially beneficial outcomes. The aim is not necessarily to restrict research or to contest any particular scientific theory but to yoke the S&T behemoth to ends chosen by the people after due deliberation and debate. ([21]: 103)

Crow and Sarewitz ask:

If we flip the current S&T policy approach on its head, and start by thinking about desired social outcomes, rather than desired inputs to the R&D enterprise (i.e., more money), where would we begin? We might identify several very general categories of outcomes that most people would agree are worth thinking about. For example:

- Social equity: the distribution of the benefits of science and technology.

- Social purpose: the actual goals of societal development that we want to pursue or advance.
- Economic and social enterprises: the shape and make-up of the institutions at the interface between technology and the human experience. ([15]: 98)

And Nordmann advises taking ‘as the starting point for ethical deliberation our historically contingent situation’ ([5]: 41), and—to recall—suggests that

Rather than adopt a believing attitude towards the future, an ethics beholden to present capabilities, needs, problems, and proposed solutions will begin with vision assessment. Envisioned technologies are viewed as incursions on the present and will be judged as to their likelihood and merit: How reliable are these claims, and do these technologies solve acknowledged problems? More generally: What do these visions tell us about the present, what is their implicit criticism of it, how and why do they require us to change? ([5]: 41)

All of these approaches propose to lodge us firmly within the field of current concerns, and suggest that, if scientists’ ideas about which future direction to take require us to abandon our current concerns, then they had better come with good reasons for us to take them seriously. In this respect, to varying degrees, these approaches more closely resemble descriptions of how most people in fact view novel projects than aspirational accounts of how we might best proceed. It may be true that visions of the future, if we are to be able to take them seriously, must offer a way of moving from our current predicament to the envisaged future one. But, it is unwise to demand—as Nordmann does—that such visions should be ‘beholden to present capabilities, needs, problems, and proposed solutions’, or—as Crow and Sarewitz do—that research and development in science should be governed by ‘desired social outcomes’. Present capabilities, needs, problems, proposed solutions, and the selection of socially desirable outcomes are often influenced by fads, bias, propaganda, prejudice, and misconception, and as such they can hardly be taken as the optimum starting point for deciding how our future should turn out. Given this, we have reason to suspect any vision of the future that is beholden to present concerns of being short-sighted: of offering a solution to current problems when taking a more

long-term view could either solve the current problems along with important problems that may arise in the future, or dispense with the need to address the current problems. Consider, for example, that an important current problem is the existence of diseases caused by harmful organisms such as bacteria. Prescribing antibiotics is a practice that is ‘beholden’ to this current problem, in that the practice has arisen as a response and solution to the problem. However, this practice has given rise to the problem of acquired antibiotic resistance: the ability of harmful organisms to become resistant to treatment by antibiotics. Whilst this is a serious problem in that it could render all antibiotic treatment ineffective, it is not a *current* problem: the diseases in question are still treatable with the right medicines. Even so, we may deem it prudent to take a long-term view by exploring options to treat diseases without using antibiotics, or to reduce acquired antibiotic resistance. Indeed, the UK government takes this view, and has begun a campaign to educate the public about how to minimise antibiotic resistance [22]. If ethicists are to be restricted to considering current and emerging technology, then allowing current social concerns to shape the direction of science is a short-sighted strategy that risks failing to emphasise the importance of taking a long-term approach. The best long-term approach sometimes requires allocating resources to projects that neither directly address any immediate need or problem, nor are known to have a high chance of success.

Conclusion

If reigning in speculative ethics is unacceptable, how else can we ensure that ethicists use their energies wisely? There is no simple answer to this question, but I hope to have shown that prioritising those issues that are currently most important or most likely is not the answer. I have suggested that disagreements about which issues are worth addressing may, after all, stem not from different attitudes to probability, but to different values—and attending to values may be the most productive way forward. Several of the examples we have discussed here, including the example of the search for a cure for cancer, demonstrate that desirable visions of the future can shape the progress of technology: if reaching a particular goal is important enough, we often judge it worthwhile to

work towards it even if the chances of success look to be unpromising. However, arguably, there is in general too little discussion of which goals are the most important ones to aim at. Governments often aim at those goals most likely to get them re-elected rather than those that will be most beneficial in the long term, and researchers often pursue those projects that are most likely to attract funding rather than those that best use their expertise for solving the most important problems. Reflecting on where our most important values lie, and how we might work to maximise them, is surely an important step towards ensuring that ethical concern, and other valuable resources, are not squandered. And, as philosophers have long known, one of the most effective ways of discovering deeply-held values involves speculating about incredible scenarios.

References

1. Parfit D (1984) *Reasons and persons*. 2nd edn. Oxford University Press, Oxford
2. Dennett DC (1978) *Brainstorms: philosophical essays on mind and psychology*. Bradford Books, New York
3. Wittgenstein L (1967) Zettel. G. Anscombe and G. von Wright (Eds) Translated by G. Anscombe. Blackwell, London
4. Urquhart A (2004) Complexity. In: Floridi L (ed) *The Blackwell guide to philosophy of computing and information*. Blackwell, Malden, pp 18–27
5. Nordmann A (2007) If and then: a critique of speculative nanoethics. *NanoEthics* 1:31–46. doi:10.1007/s11569-007-0007-6
6. O’Neill O (1986) The power of example. *Philosophy* 61 (235):5–29
7. Manson N (2002) Formulating the precautionary principle. *Environ Ethics* 24:263–274
8. Keiper A (2007) Nanoethics as a discipline? *New Atlantis (Spring)*:55–67
9. Levi WL (1943) Temperament and moral theory. *Ethics* 53 (2):128–132. doi:10.1086/290336
10. Sunstein CR (2004). Your money or your life. *The New Republic*, 11 March, 2004. Retrieved 8 August, 2008, from http://www.powells.com/review/2004_03_11
11. Joy B (2000) Why the future doesn’t need us. *Wired* 8 (04):238–262
12. Brown JS, Duguid P (2001) Don’t count society out: a reply to Bill Joy. In: Denning PJ (ed) *The invisible future*. McGraw-Hill, New York, NY, pp 117–144 Retrieved 13 August, 2008, from http://people.ischool.berkeley.edu/~duguid/SLOFI/Don't_Count.htm
13. Slovic P, Finucane M, Peters E, MacGregor DG (2002) The affect heuristic. In: Gilovich T, Griffin D, Kahneman D (eds) *Heuristics and biases: the psychology of intuitive judgement*. Cambridge University Press, New York, NY, pp 397–420

14. Starr C (1969) Social benefit versus technological risk. *Science* 165:1232–1238. doi:10.1126/science.165.3899.1232
15. Crow MM, Sarewitz D (2001) Nanotechnology and societal transformation. In: Teich AH, Nelson SD, McEnaney C, Lita SJ (eds) AAAS Science and Technology policy yearbook. American Association for the Advancement of Science, Washington, D.C, pp 89–101
16. Weckert J, Moor J (2007) The precautionary principle in nanotechnology. In: Allhoff F, Lin P, Moor J, Weckert J (eds) *Nanoethics: the ethical and social implications of nanotechnology*. Wiley-Interscience: Hoboken, NJ, pp 133–146
17. Shaw GB (1921) *Back to Methuselah*. Penguin, London
18. Tversky A, Kahneman D (1973) Availability: a heuristic for judging frequency and probability. *Cognit Psychol* 5:207–232. doi:10.1016/0010-0285(73)90033-9
19. Boyle A (2008). Doomsday fears spark lawsuit. *Cosmic Log*, 27 March, 2008. Retrieved 2 July, 2008, from <http://cosmiclog.msnbc.msn.com/archive/2008/03/27/823924.aspx>
20. Bostrom N (2006) Welcome to world of exponential change. In: Miller P, Wilsdon J (eds) *Better humans? The politics of human enhancement and life extension*. DEMOS, London, pp 40–50
21. Bostrom N (2007) Technological revolutions and the problem of prediction. In: Allhoff F, Lin P, Moor J, Weckert J (eds) *Nanoethics: the ethical and social implications of nanotechnology*. Wiley-Interscience, Hoboken, NJ, pp 101–118
22. Department of Health Antibiotic Campaign 2007–2008. Retrieved 24 July, 2008, from http://www.dh.gov.uk/en/PublicHealth/Patientsafety/Antibioticresistance/DH_082512