

Maximising Happiness

Jonathan Bennett

from: R. I. Sikora and B. Barry (eds), *Obligations to Future Generations* (Temple University Press, 1978), pp. 61–73.

1. When it is wrong to bring into existence someone who will be miserable, what makes it wrong is not the threat of misery hanging over the possible person, but rather the fact that if one does it there will be real misery for an actual person. This belongs in the same category as the wrongness of making a happy person miserable, or of failing to make a person less miserable than he is. These are all matters of the (dis)utilities—the ill-fare and welfare—of present and future actual people.

One might assign (dis)utilities to past people, e.g. objecting morally to using the calculus for military purposes because Leibniz wanted all his discoveries to contribute to universal peace—‘If I use the calculus in building a bomb, I am bringing a disutility to Leibniz by bringing it about that he was to that extent a man whose hopes were not going to be realized’. (It seems that this would have been acceptable to Aristotle, who said that someone’s ‘happiness’—according to the standard translation—may be affected by what happens after his death.¹) I am not endorsing this attitude to past people, but I shan’t quarrel with it here. For my present purposes, I allow it into my scheme because the people in question are at some time actual.

But no sense attaches to ‘(dis)utility for a person who is at no time actual’. So if a failure to bring someone into existence is ever wrong for utilitarian reasons, these must concern the utilities of people who are at some time actual, not those of the person whose coming-into-existence didn’t happen. It might be wrong for me to fail to beget a child because that would deprive my parents of the pleasures of grandparenthood, or because any child of mine would be sure to benefit mankind; in one case my parents are deprived, in the other mankind in general. But it couldn’t be wrong because by not bringing the child into existence one deprives *it* of something.

2. ‘But it could be wrong because by bringing the child into existence one gives it something.’

Well, I submit that in so far as moral questions are to be determined by facts about personal utilities, it should be through the rule:

The question of whether action A is morally obligatory depends only upon the utilities of people who would exist if A were not performed,²

taking this to include people who did or do or will exist whether or not A is performed. Nothing is excluded except

¹ Aristotle, *The Nichomachean Ethics*, Bk. I, ch. 10.

² In this, as indeed throughout this paper, I follow the excellent lead of Jan Narveson, ‘Utilitarianism and New Generations’, *Mind* 76 (1967), pp. 62–72.

the utilities of people who would not exist if A were not performed—those who would exist if and only if A were performed, and also, trivially, those who would not exist whether or not A were performed.

This rule does not express anything that could be called ‘negative utilitarianism’¹ It is not part of any general moral scheme according to which one may be obliged to limit harm but never to bring help, or anything of that kind. Nor does it assume that morality basically tells one what not to do, but doesn’t tell one what to do: the difference between ‘saying what we should do’ and ‘saying what we shouldn’t do’ is just the difference between specific instructions and highly general ones; and nothing like that is involved in the above rule about utilitarian reasons for moral obligations. The rule is ‘negative’ only in the sense of implying that when I am considering whether I am obliged to do A, the only people whose utilities I need to consider are those who do or would exist if I did *not* do A.

The obvious alternative to my rule is the more generous one that says:

The question of whether action A is morally obligatory depends upon the utilities of people who would exist if A were not performed, and of people who would exist if A were performed.

Both rules allow that one may be obliged not to bring someone into existence because he is bound to be wretched. What about the obligation to bring someone into existence because he is bound to be happy? The longer rule—the one that makes provision for the larger array of obligations—permits such a reason for such an obligation, whereas my shorter rule forbids it.

Over most of the territory, the two are equivalent. They differ only when there is a question about bringing people into existence: then the short rule confines us to moral injunctions of the form ‘Don’t do it unless. . .’, whereas the long one leaves the door open for further injunctions of the form ‘Do it if. . .’. The main reason for preferring the short rule is that it is harmless while the long one opens a Pandora’s box of troubles. The long rule doesn’t actually force us into trouble: since it merely *permits* us to have a certain sort of moral principle we don’t break or flout it if we decline the invitation and keep our morality within the limits laid down by the shorter rule. But attempts have been made to accept the long rule’s invitation, i.e. to devise moral principles that are forbidden by the shorter rule—principles that enjoin the creation of happy people, roughly speaking—and so far every one of these attempts has had morally intolerable consequences.

Unless some powerful positive reason can be found for the long rule. therefore, i.e. some reason for wanting a stronger set of principles of obligation than the short rule permits, the case for the short rule will be overwhelming.

3. Some philosophers seem to have thought that they had a case for the long rule. i.e. for the ‘Do it if. . .’ clause as well as for the ‘Don’t do it unless. . .’ clause, They seem to have relied on the idea that if doing A would bring into existence someone who would then be happy, then not doing A does involve *a loss of utility*—namely the loss of that happiness—even though no actual person is deprived of the utility.

That in its turn appears to depend upon the notion, which I believe to be noxious in utilitarian morality, of *amounts* of. . . The gap can be filled by whatever form of utility you

¹ As is suggested, against Narveson, by Jonathan Glover, *Causing Death and Saving Lives* (Penguin Books, 1977), p. 67.

like: I shall say ‘amount of happiness’, but my concern is with ‘amount’, not ‘happiness’.

Now, it doesn’t matter whether it is a total amount, an average amount, or some compromise between these: the mere notion of *amount* lets philosophers introduce a surrogate for the proper notion of utility—it gives them utilities that are not *someone’s*, in the form of quanta of happiness that nobody has but that somebody could have. As well as deploring the situation where a person lacks happiness, these philosophers also deplore the situation where some happiness lacks a person; and they speak of the latter situation as being one in which some utility is lost. This, I submit, is a philosophical mistake.

It’s the mistake of inferring from ‘We ought to make people as happy as possible’ that ‘We ought to produce as much happiness as possible’. It doesn’t follow.

I don’t say that it is a *philosophical* error to hold, as a basic item in one’s moral code, that one ought to produce as much happiness as possible. The philosophical mistake is to think that this—or the part of it that doesn’t concern actual people—has to do with ‘loss of utility’ in some reasonable sense of that phrase.

4. With the muddle about ‘loss of utility’ set aside, perhaps there is some other reason to think that one ought to try to maximize happiness—taking this, loosely, to mean that in some circumstances one would be morally obliged to bring people into existence because they would then be happy. Might one not, for instance, base this on the view that the more happiness there is the better?

Suppose that we have believed that a certain region of the universe is cold and sterile, and then we discover that it is teeming with life and that its sentient inhabitants have very happy lives. Isn’t this good news?

In my value-system, Yes; but not because I think that the more happiness there is the better. I share Leibniz’s liking for rich, organic complexity, and so the discovery that our world has more of it than we had realized would be good news indeed. But amounts of happiness don’t come into it. Happiness is relevant only in that the extra organic complexity would not be very welcome if the organisms were desperately miserable. From the fact that I am glad they exist, and wouldn’t be glad that they exist if they weren’t also happy, it doesn’t follow that my gladness stems from a wish for maximal amounts of happiness.

Still, there could be someone whose gladness did have that source—someone who got a lift out of the sheer thought of all that extra happiness. That attitude—which might be expressed as the view that happiness is a good, or that the more happiness there is the better—seems intelligible as a good-hearted response to the thought of happiness in others. But it starts to slip out of my grasp, and to become something that I can only take to be an outcome of philosophical confusion, when it is thought to imply that one is morally obliged to bring people into existence so that they can be happy. I’ll return to this point in §9 and §10 below.

5. This issue has a bearing on the question of whether *Homo sapiens* should be allowed to continue. There is obviously a moral case for letting our species die out: that would benefit many other terrestrial species, it would prevent the misery that will be suffered by millions of people if the species does continue, and so on. Against that, there are moral reasons for keeping the species going: reasons stemming from the disutilities to actual people of the phasing out of *Homo sapiens*—the frustration of the hopes of would-be parents, the loneliness of the final few thousand, and so on. From now on, I shall set all of that aside.

If we knew that all of the foregoing reasons pretty much balanced out, so that out of consideration of the (dis)utilities of people who are at some time actual we couldn't make a strong case for or against giving *Homo sapiens* a future, would that be the end of any moral consideration of the question?

Well, I am committed to saying that the question shouldn't be kept alive through any such thought as: We ought to perpetuate our species because if we do larger amounts of happiness will be had than if we don't. But I don't have to say that no other moral considerations could be brought to bear.

For example, someone might accept a principle enjoining the preservation of every species, or every animal species, or every instance of extreme physical complexity, or every form of life that is capable of moral reflection, or. . . Any one of these could be brought to bear on the 'future for mankind?' question, in addition to all the 'utilitarian' considerations properly so-called; and clearly none of them enjoins the maximization of happiness or anything like it.

Still, I don't want to make much of that. Although none of those principles is any worse than the principle that we ought to maximize happiness, none of them is much better either: none of them seems to have much to recommend it, considered as an item of basic morality. So I still have not identified any acceptable moral ground for an obligation to perpetuate our species, other than grounds involving the utilities of actual people.

6. When the utilitarian pros and cons are set aside, someone might be simply indifferent about the perpetuation of our species. This attitude, though it is perhaps unusual—even unnatural—should not be found shocking or ugly.

It is not my attitude, however. I am passionately in favour

of mankind's having a long future, and not just because of the utilities of creatures who were, are, or will be actual. This is a practical attitude of mine for which I have no basis in general principles. The continuation of *Homo sapiens*—if this can be managed at not too great a cost, especially to members of *Homo sapiens*—is something for which I have a strong, personal, unprincipled preference. I just think it would be a great shame—a pity, too bad—if this great biological and spiritual adventure didn't continue: it has a marvelous past, and I hate the thought of its not having an exciting future. This attitude of mine is rather like my attitude to pure mathematics and music and philosophy: even if they didn't have their great utility, I would want them to continue just because they are great long adventures that it would be a shame to have broken off short.

In saying that this is not a principled attitude, I am partly saying that it is not *based on* any principles. If someone doesn't share this attitude of mine, there is little that I can say to him on the topic: I can't reach down for something that produces or confirms these attitudes in me, and appeal to the dissentient at that level.

Some practical attitudes are not based on principles because they *are* principles, i.e. are basic in the person's morality. But my wanting a future for humanity is not like that either. Indeed, it hasn't the right logical shape to constitute a universal principle: it is not a stand in favour of there being animals that answer to this or that general biological or psychological description, but rather a stand in favour of there continuing to be animals descended from some of *these*—and here I gesture broadly towards all my human contemporaries. The essential reference to particulars sharply differentiates this from such principled attitudes as my opposition to the causing of unnecessary pain to any creatures that can feel pain.

7. It is because of the unprincipled nature of my pro-humanity stand (as I call it, for short) that I don't regard the stand as part of my morality or, therefore, as a source of moral obligations. Hare's thesis that only universalisable practical attitudes should be accounted moral is true, I believe, and I find it helpful and illuminating. But for present purposes I needn't insist upon it. If you think there can be unprincipled moral stands, then you may count my pro-humanity stand as 'moral' after all; but you won't demand to know what principle(s) it involves, and that is what matters.

It matters because it frees me from the risk of being committed to absurdities. My pro-humanity stance, not resting on any principle, isn't vulnerable to the production of embarrassing counter-examples. You can describe to me a possible world, and ask whether I favour the continuation of *Homo sapiens* in that world (always setting aside the genuinely moral considerations, e.g. utilitarian ones); and I shall answer on the basis of how far it strikes me as a pity, a shame, too bad, if the story were to be cut short in that world. There might be patterns and regularities in my answers; you might even come to be able to predict them. But I wouldn't be applying a principle, and so I would never suffer the embarrassment of being committed to saying one thing while wanting to say something different.

Although my pro-humanity stance is not based on a principle, it does have some inner structure and I do have something like reasons for it. My attitude to mankind's future is conditioned by my attitude to its past: my sense that it would be a shame if the story stopped soon is nourished by my sense that it has been an exciting story that involves some long-term endeavours that aren't yet complete. I would probably care less about the 21st century if I didn't love the 17th so much. Thus, if I consider a world where there is no

great disutility to actual creatures if *Homo sapiens* dies out, and where the continuity couldn't be more than a purely biological one because all the extrasomatic information and most of the non-genetic information have been irretrievably lost, I have no sense that it would be a great shame if the thread were to be broken then and there, and the species brought to an end.

So if I *had* to find a principle that I could slide in under my pro-humanity stance, it would probably be one about the prima facie obligation to ensure that important business is not left unfinished. Once that principle was clearly formulated, the counter-exemplifying game could begin. But perhaps I could so formulate my principle that no embarrassing counter-examples could be found. Anyway, I would be free from the peculiarly horrid difficulties that beset theories relying on the notion of amounts of possibly unowned happiness.

8. Although I know roughly what principle I would have if I had any, and although I doubt if it would lead me into grave and demonstrable trouble, I don't in fact have a principle. I say this partly because, as I said earlier, my attitude towards our great adventure is not an attitude I have towards all great adventures, nor even towards all of some general kind. But as well as that point about the logical structure of my attitude, there is also phenomenological evidence that the attitude is not a principled, moral one. It concerns the way I react to the thought of dissent—including my own possible future dissent—from my present attitude towards the continuation of mankind. I shall explain this.

When we accept something as a matter of moral principle, we engage ourselves to treat it somewhat like a matter of objective fact. I don't think that there are any moral facts, but I still describe moral views as 'correct' or 'incorrect', and

I think of moral disagreement as implying that someone is in error. Thus, if I imagine myself coming to relinquish a basic moral view that I now hold, I am imagining a deterioration in my moral sensibility. In contrast, when I imagine myself having a different attitude to mankind's future (still, as always, setting aside the utilities of actual people), I don't see myself as *spoiled* but just as *changed*.

There is an analogous contrast between my attitudes to someone else who disagrees with me on a basic moral matter and to someone who doesn't share my attitude to the future of humanity. With the latter, I might try to bring him around by capturing his imagination with a certain picture of the human adventure; but if I failed, it would not induce the sort of distress that is appropriate when one is faced with fundamental moral disagreement.

If the question of mankind's future became a matter of practical urgency, and the utilities of actual people left the issue evenly balanced (!), then I might have practical reasons for wanting to bring dissentients around to my way of seeing things. That would be a simple matter of wanting them to want what I want, so that they would act to produce what I want: this, I submit, is very different from the sort of pressure towards agreement that obtains in a matter that is clearly classified as moral.

Those remarks imply something that I had better make explicit. Although my pro-humanity stance is not principled, it is thoroughly practical, and it may be at least as strong as some of my principled, moral attitudes. Don't play it down by calling it 'aesthetic' and letting that label suggest that it is merely contemplative. If the occasion arose, I would work for it, and probably even suffer for it.

I would also fight for it. If you try to bring it about that there is no 21st century human species, linked by me (and my contemporaries) to Spinoza and Aristotle, you may not be acting wrongly but I'll fight you all the same.

If my attitude in this matter is just an unprincipled desire, would I be morally in the right if I fought in defence of it, i.e. if I tried to make this desire of mine prevail over the desires of others? I think so. I don't mean that it is morally all right to fight anyone, in any manner, in the furtherance of any desire. But I do hold that each person is morally entitled to give some special weight to his own wants and needs and interests, just *qua* his. Although this 'morality of self-interest', as I call it, seems to be neglected in philosophical writings about morality, it looms very large in almost everybody's moral thinking. I have no hesitation in endorsing it.¹

9. I suspect that my sort of attitude to the future of mankind is very widespread. It goes naturally with how theists, when not in the toils of the problem of evil, think about God as creator. I'll put this point in a non-theological manner, just for comfort. Suppose that you have a chance to bring into existence a race of creatures whose lives will be tremendously worth living. If you take the trouble to create them, they may well bless you for it, praise you, thank you, rejoice in the very thought of you—all the things that go with their seeing your act as one of supererogation rather than as something you were obliged to do. Why then should you see it any differently?

(I moved from God to you because of general difficulties about moral attitudes towards God—e.g. why lavish praise and gratitude on Him for something that was, for Him, no trouble at all?)

¹ (Added in 2012:) I discuss it at more length in my *The Act Itself*, Oxford 1995, sections 48, 50.

For another slant on the matter, suppose that we came to share this planet with another species who were superior to us in every way except that we still had control (e.g. the Others were vulnerable to radio-waves, and had no comparable means of controlling us). The question might arise as to whether we should hand over dominance to these Others, and, since we could never happily accept a subordinate role, allow *Homo sapiens* to die out. Of course many genuinely moral considerations would come in, especially about the utilities of actual humans and actual Others; but what about the residual question—the question about the sheer idea of humanity’s abdicating from its own future?

The example is adapted from Isaac Asimov, whose attitude to the prospect is an enthusiastic *What a way to go!* Many thoughtful people, I believe, would share his feeling that there is something splendid about the idea of a biological species voluntarily ceding its place to another which it judged intellectually and morally superior. And many would feel, instead or as well, that there is something profoundly sad about it. I wonder how many people, if they hadn’t been misled by bad philosophy, would want to add to those two civilized responses a moral judgment about an obligation to maximize happiness?

10. By willing the continuation of *Homo sapiens* one is inevitably willing profound misery for many people who would escape it if the species were allowed to die out. It has been pointed out that if that is not to be downright immoral, there must be something one can set against all that misery—something that outweighs it in the moral scales.¹ That is correct, but it falls triply short of re-establishing the amount-of-happiness approach, as I now show.

(a) It is plausible to think that the whole case for the continuation of *Homo sapiens* rests on the utilities of actual people—i.e. ones who were, are or will be actual even if mankind dies out. These utilities are considerable: the phasing out of our species over 150 years, say, would involve ghastly horrors; and the disutilities also include events which, though not intrinsically dreadful, serve to frustrate hopes, abort plans, disappoint expectations, and so on (as when a happy man suffers a disutility by dying quietly in his sleep). Perhaps, if it were not for all that, it would be morally inexcusable to sanction the continuation of mankind.

That would probably be the view of anyone who held, as many do, that if we were freely created by a God who knew what was coming, then what He did was inexcusable. Such a view does not commit one to perpetual gloomy discontent: someone who is radiantly satisfied with his own lot might still sincerely judge that the price was, morally speaking, too high. Nor is it implied that one should work for a state of the world in which mankind could be phased out without too much disutility to actual people: improved birth control techniques, education aimed at discouraging long-term plans and hopes, help for the morale of the gradually shrinking remainder, and so on. Any energies devoted to tackling the problem from that angle could as well be spent on trying to reduce the likelihood that millions will be wretched if *Homo sapiens* does have a long future.

(b) Someone who thinks that there is a moral case for the continuation of mankind, over and above one stemming from the (dis)utilities of actual people, need not immediately fling himself into the arms of a principle enjoining the maximization of happiness. There are other bases he might appeal to, some of them more plausible than any amount-of-happiness

¹ R. I. Sikora, ‘Utilitarianism: The Classical Principle and the Average Principle’, *Canadian Journal of Philosophy* 5 (1975), pp. 409–419, at pp. 414–415.

principle: e.g. the principle that enjoins the completing of unfinished business, mentioned at the end of § 7 above.

(c) Someone who isn't prepared to adopt either of those positions can still escape the conclusion that we are morally obliged to maximize happiness. I shall state this in an abstract form. Suppose someone says something of the form: 'Although if *Homo sapiens* continues many people will be unhappy, this is outweighed—the continuation of the species is made morally permissible—by the fact that if the species is continued then P', where P may be 'many people will enjoy great happiness', or 'certain intellectual and spiritual endeavours will continue', or anything else that doesn't involve the utilities of people who are or will be actual even if mankind does not continue. My point concerns not the identity of P, but rather its formal role in the argument. It is here being adduced as a reason for something's being morally permissible, but does that necessarily imply that it can also help to make something morally obligatory? Unless an affirmative answer to that can be defended, the argument falls at this fence even if it has surmounted the previous two.

The best case I can think of for an affirmative answer goes like this. 'If there is a *prima facie* case against doing A, one is permitted to do it only if it will bring enough compensating good; and if doing A would do a large enough amount of good, as compared with what can be brought about through some (any) alternative to A, then one is downright obliged to do A. Thus, the raw materials for a permission can always be made to serve to generate an obligation.'

Although I cannot disprove that, it is not compellingly self-evident, and I for one don't accept it. In particular, I believe that it fails in cases where the 'goods' in question are personal utilities and the practical issue affects what people will come into existence. I now explain this.

In § 2 above I presented and recommended a rule that says that facts about personal utilities are relevant to moral obligations only through arguments of the form: 'It is obligatory to do A because if A is not done then there will be people who...'. That lies at the heart of my rejection of the idea that we are obliged to maximize happiness even where this involves creating people to have the happiness.

I stand by that rule because I see no reason not to. So far as the present difficulty is concerned, all I need is to show that the rule allows me to argue: 'The continuation of mankind is morally permissible, despite the predictable misery that will be involved, because of all the happiness that will be involved.' I now show that the rule does indeed present no obstacle to such an argument.

Since we are to relate a rule about obligation to a question about permissibility, the latter should be re-phrased in terms of the former. Our question, then, concerns the *obligatoriness of the NON-continuation of mankind*: the question of whether it is permissible for mankind to be given a future is the question of whether it is obligatory for mankind to be denied a future. For short, let 'A' stand for something like 'bringing it about that mankind has no future', Then our question is: *Is A morally obligatory?*

Now, my rule in §2 lets us bring to bear upon that question any facts about the utilities of people who would exist if A were not done, i.e. if mankind were allowed to have a future. As well as making a negative case out of the fact that if A is not done there will be many miserable people, the rule lets us make a positive case out of the fact that if A is not done there will be many happy people. I don't say that it is all right thus to balance the bliss of some against the misery of others; but if it isn't, *cadit quaestio*. What I am saying is that my rule gives me access, in considering the obligatoriness of the non-continuance of mankind, to all the facts about all

the people who will exist if mankind's non-continuance does not happen, i.e. who will exist if mankind continues. And so I have a clear, coherent basis upon which to maintain that the proposition *If mankind has a future then there will*

be many happy people has some tendency to make it morally permissible to give or allow mankind a future, while having no tendency whatsoever to make it morally obligatory to give or allow mankind a future.