Soc Choice Welfare (1991) 8: 79–88

Social Choice and Welfare

Should we be very cautious or extremely cautious on measures that may involve our destruction?

On the finiteness of our expected welfare*

Y.-K. Ng

Monash University, Clayton, Victoria, Australia 3168

Received November 15, 1989/Accepted October 3, 1990

Abstract. For decisions (e.g. nuclear power development, environmental protection, genetic engineering) that may affect the probabilities of the continued survival of the human race, whether we should be very cautious or extremely cautious (defined as refusing to undertake anything that may reduce our survival probabilities) depends on whether our expected welfare is finite or infinite. If it is infinite, a paradox arises in the trade-off between our own expected welfare and that of future generations, since a small fraction (probability change) of infinity is still infinite. However, limitations on population size and average welfare suggest a finite expected welfare but the possibility of transforming our own selves perhaps by genetic engineering increases our expected welfare tremendously but still finite.

I. Introduction

Many important public policies affect the long-term welfare (and even survival) of the whole world, e.g. nuclear power, environmental viability, genetic engineering. For decisions involving such grand issues, whether we should be conservative or adventurous depends crucially on our assessment of our expected welfare, i.e. our welfare through to infinity discounted by the uncertainty of our continued survival. For example, undertaking something (e.g. building nuclear power stations) with a certain gain may also involve a certain risk of reducing our survival probability. Whether the gain is worth the risk critically depends on the loss we place on non-survival or the value we place on continued survival. The latter depends on our expected welfare.[1] In particular, if we assess our expected welfare to be infinitely large, we should rationally be extremely con-

* The author acknowledges the helpful comments of John Broome and Maurice Salles.

[1] Personally, I regard it as not only dependent on but in fact equal to our expected welfare. Being a welfarist, I place no further value except welfare. However, for the purpose of this paper, readers may only have to agree that welfare is an, even if not the only, important issue.

servative, refusing to undertake anything that may threaten our survival unless the gain is also infinite.[2]

It is true that, in the real world, most decisions are undertaken by nations and individuals which have finite life expectations. Thus there are issues of externalities of the action of one nation on those of others, and of the present generation on those of the future generations. However, to concentrate on the problem in its pure form, we shall abstract from such complications and assume that we (i.e. human beings) act as a unified body maximizing our expected welfare through to infinity (except Sect. II where we discuss the contrast between the present and future generations). For simplicity, we also ignore the welfare of non-human sentients or hold it constant.

The next section examines the conflict between the present and future generations. In particular, it is pointed out that, if the expected welfare of mankind is believed to be infinite, a paradox or dilemma exists in our trade-off between our own expected welfare and that of future generations. With some but imperfect concern for the welfare of future generations, we would like to weight our welfare more than that of the future generations. However, we may not be able to put this weighting bias into practice since we should not reduce the survival probability of future generations despite the weighting bias because a small fraction (probability change) of infinity is still infinite. In Sect. III, the finiteness of our expected welfare is examined. It is argued that, due to the limitations on the number of people the universe can sustain and the biological limitations on the welfare level of each of us, our expected welfare should be finite. Section IV examines the possibility of raising the welfare level through the transformation of ourselves perhaps by genetic engineering. Despite this possibility, it is still argued that our expected welfare is finite.

II. Our self interest versus the welfare of future people

While we as a species may live for thousands of years, none of us living now will survive for much more than a hundred years if at all. Thus, there is the problem of a possible conflict between the present and future generations. It is true that the conflict is not total since those living in the present generation do care about the welfare of future generations, especially their immediate offsprings, as analysed in the interlocking generations models. However, it is unlikely that the concern will be complete in the sense that we of the present generation put the welfare of those in the future generations on a par with our own welfare such as implied by the maximization of the expected welfare of the present and future generations combined.

Given incomplete concern, the present generation will put more weight on its own welfare. It will then be less cautions or more willing to undertake measures that increase its own welfare at the risk of destroying the future. An interesting dilemma arises if the present generation believes that the expected welfare of the future generation is infinite. No problem arises if the present generation is perfectly self-concerned and hence maximizes its own welfare irrespective of the costs (including the risk of destruction) imposed on the future. If the present generation

[2] On the rationality of maximizing expected welfare in the presence of uncertainty, see Ng (1984).

is not perfectly selfconcerned and is, for example, just willing to trade off one unit of its own welfare for N units ($1 < N < \infty$) of future welfare, a dilemma arises as outlined below. (For simplicity, we ignore here the distinction of future generations closer to the present and further into the future and ignore individuals and groups within the present or the future generations. Thus, we concentrate on the contrast between the present generation and the future.)

If the expected welfare of the future generations is infinite, even a very small risk of destroying this future involves an infinite cost. This infinite cost, even discounted by multiplying with $\alpha (= 1/N)$, is still infinite. Then, even if the present generation is not perfectly moral (perfect morality defined by $\alpha = N = 1$; imperfect morality by $0 < \alpha < 1$), it seems that it has to choose as if perfectly moral for measures involving risks of destroying the future. How could the present generation put its bias towards its own welfare into practice?

One possibility is to trade-off units of present welfare not with units of future welfare but with the percentage probability of future destruction. For example, we (i.e. the present generation) may be (just) willing to accept a 0.1% probability of destroying the future if our welfare could be increased by x units. This may seem a reasonable way of trading-off present and future welfare. However, when we take account of the trade-off between units of present and future welfare with certainty, some form of inconsistency can be shown to be involved.

When faced with mutually exclusive alternatives of increasing our own welfare by a certain amount with certainty versus increasing future welfare by another (presumably larger) amount with certainty, our choices may imply our willingness just to trade-off one unit of our welfare with N units of future welfare, with $1 < N < \infty$. However, our willingness to accept a 0.1% probability of destroying the future if our welfare could be increased by x units ($0 < x < \infty$) implies that we are willing to sacrifice an infinite number of units of future expected welfare for a finite gain in our own welfare. Clearly, both the present and future generations could be made better off (higher expected welfare) by some change in our trade-offs in increasing N (which makes us better off and the future worse off) and increasing x (which makes us worse off and the future better off). Such a combined change that makes both the present and the future better off is always feasible as long as N and x are not infinite. (This is so since the expected welfare of the future generations can be increased by a small reduction in the risk of destruction even at a big increase in N sufficient to more than compensate the present generation.)

At the risk of repetition, it should be emphasized that the above paradox or dilemma applies not just to a believer in utilitarianism or welfarism. Even for an "extra-welfarist" (who believes in the ultimate importance of things other than welfare; such beliefs are usually due to the confusion of non-ultimate considerations with basic values, as argued in Ng, forthcoming), the paradox still applies, if the relevant non-welfarist factors are held unchanged or if their changes are not significant enough to offset the infinite reduction in the expected welfare of the future generations.

It might be argued that, if the expected welfare of the future is infinite, any finite change is not going to matter since infinity minus any finite amount is still infinity. We should thus make N infinite even if we do care for the future. This argument holds if each generation in the future will enjoy infinite welfare. But this is clearly not the case. We might get infinite expected future welfare only by integrating a finite level of welfare infinitely into the future. If this finite level

increases faster than the uncertainty discount, the integral becomes infinitely large. If we can increase say the welfare of the next three generations by a sufficiently large amount at a sufficiently small cost to us (the present generation) and no effect on the welfare of people three generations after, it is clearly a good bargain if we have any concern for future welfare, even if we believe that expected future welfare is infinite.

It may seem that one way out of the difficulty is to adopt a lexicographic ordering, maximizing present welfare but preferring higher expected future welfare if present welfare remains constant. However, I believe that it is irrational to have lexicographic orderings. *Either* we have some non-zero concern for future welfare in which case α is strictly positive and we are willing to sacrifice some of our own welfare if the increase in future welfare is big enough, *or* we have no concern and should be indifferent as long as our own welfare remains unchanged, irrespective of the level of future welfare. Lexicographic ordering may be an acceptable *approximation* for preferences between two (or more) aspects of enormous difference in their relative importance, but is not strictly speaking reasonable. The argument for lexicographic preferences is mainly based on such a question as, "Is there some number of trinkets that will induce a starving coolie to part with one bowl of rice?" (Chipman 1960, p. 221). The answer may well be negative, but this does not make his preference lexicographic. With divisibility, his preference can only be lexicographic if, given that he prefers more trinkets to less, there is no number of trinkets that will induce him to part with 0.00000...1 grain of rice. This is clearly unlikely. [See Ng (1983a, p. 28) on the implausibility of lexicographic preferences.]

It thus seems that unless future expected welfare is not infinite, we have a moral dilemma if we also have an imperfect concern for future welfare (almost certainly so).

III. Is our expected welfare infinite?

It may be argued that, while an individual is mortal, the human society is immortal. Assuming a positive level of welfare at any moment in time, our future welfare is thus infinite. This reasoning is wrong because we cannot be certain of our immortality. Our solar system will be running out of heat in finite time. Even if we can colonize planets in other solar systems, the whole universe itself may also be finite. If so, then the second law of thermodynamics ensures us a finite life. But is our universe a closed system? No one can give a definite answer. [For an account of the current views on the universe for laypersons, see Hawking (1988).] At least we have to accept the possibility of our mortality, perhaps well before the end of our solar system. Events such as fatal celestial collision cannot be ruled out completely. Thus, instead of infinite survival, our expectation should be something as illustrated in Fig. 1.

One possibility is an abrupt ending at some time in the future as illustrated by the curve ABC. Of course, we are unlikely to have uniform improvement over the segment AB; some ups and downs are possible. However, since we do not know the precise shape and timing of these ups and downs, the *expectation* could be a fairly smooth curve, with increasing/constant/decreasing rates of improvement as the case may be. It is also possible that the ending is not perfectly abrupt

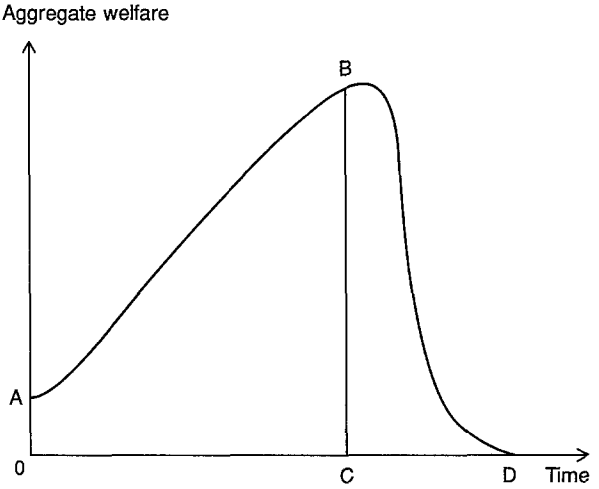

Fig. 1

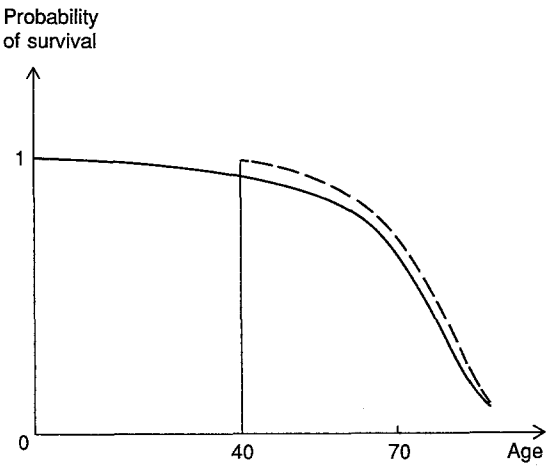

Fig. 2

and we may have a case as depicted by the curve ABD. For simplicity of exposition, let us concentrate on the case ABC with an abrupt ending.

While we expect the ending to come some time in the future, we do not know when it will occur. This is somewhat similar to the life expectancy of an individual, though in this case the probability of ending one's life around 65–85 years old is very high. In terms of the probability of survival (until at least the respective ages), an individual at birth typically faces a function approximately illustrated by the solid curve in Fig. 2. After she has actually survived to the respective ages, she should adjust her survival function accordingly. For example, at age 40 her survival function typically looks like the dotted curve in Fig. 2.

For the survival function of mankind, a lapse of 40 years in itself will cause virtually no change in the survival function. However, if our knowledge about the future prospect changes during the 40 years, a significant change in the

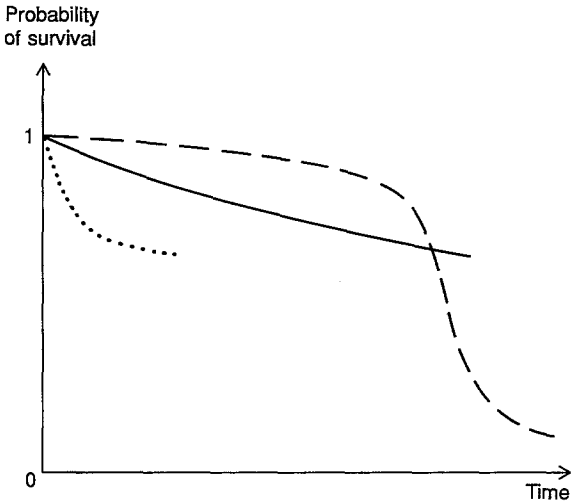

Fig. 3

function may result. If we do not believe that we will be particularly vulnerable at certain periods, our (current) survival function may look something like the solid curve in Fig. 3, with a uniform rate of decrease through time. What is the appropriate rate depends on how optimistic we view our future. If we believe we are particularly vulnerable in the next thousand years or so but will be fairly safe if we manage to survive this dangerous period, the survival function looks like the dotted curve. On the other hand, if we believe that we will be pretty safe until our sun runs out when we will almost certainly perish, our survival curve looks like the dashed one. However, there are no strong enough reasons to support either of these two contrasting views. We may be vulnerable to nuclear wars now but may also be vulnerable to something much more destructive than nuclear wars after a thousand years. The running out of the sun may be very threatening to our survival, but there will be many other threatening possibilities well before that. Moreover the timing of the sun running out is also subject to uncertainty. Furthermore, by that time we might have found ways to replace the sun or might have migrated to another planet of a younger sun. It seems that our knowledge about our future survival is so uncertain that a uniform rate of decrease is a more appropriate expectation. If this is the case, it leads to a tremendous simplification as a constant rate of discount is so much more manageable mathematically than a changing one for dynamic optimization over time.

We may define our "half life" as the length of time our probability of surviving at least that long is 50%. This is different from life expectancy which is the average expected lifespan. Let the probability of our survival depreciate continuously at a constant rate of r per annum. Then our (i.e. human) life expectancy equals

$$\int_0^\infty \mathrm{tre}^{-rt}\,\mathrm{d}t$$

$$= \left| -\mathrm{e}^{-rt}(t+1/r) \right|_0^\infty = 1/r \ .$$

In other words, for the case of a constant rate of decrease, our life expectancy is just the inverse of this rate. If we face a risk of extinction of 0.1% per annum, our life expectancy is 1,000 years. On the other hand, denoting our "half life" as x, we have

$$e^{-rx} = 1/2$$

$$x = \ln 2/r = E \ln 2 = 0.6931472E \ ,$$

where E = life expectancy. Thus for the case of a constant risk, our half-life is about 69.3% our life expectancy. With a per annum risk of 0.1%, our life expectancy is 1,000 years and our half-life is 693.15 years. This can be confirmed by a simple calculation.

Figure 4 shows how our future can be illustrated compactly in terms of our half-life. The horizontal axis is scaled as indicated.

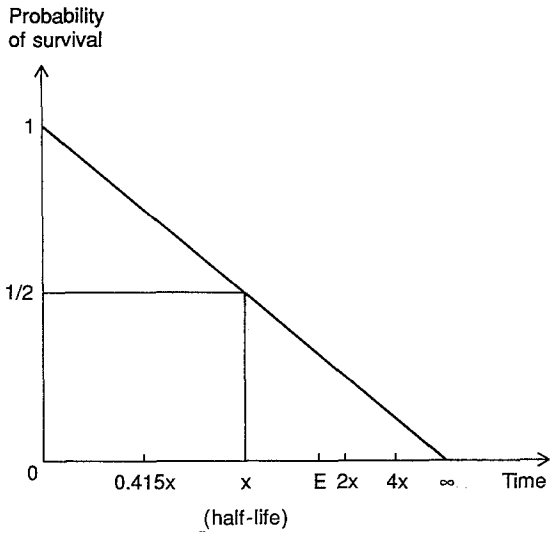

Fig. 4

With any given constant rate of decrease in our survival probability (and hence a constant rate of uncertainty discount), our expected welfare may be infinite if our instantaneous welfare can be expected to increase exponentially at a rate higher than the rate of decrease in our survival probability. This is so since our expected welfare is simply our instantaneous welfare times the survival probability and integrated from now to infinity.

$$EW = \int_0^\infty P_t W_t \, dt \ ,$$

where EW = expected welfare, P = probability of survival until time t (viewed from time 0), W_t = instantaneous welfare at time t. If P_t decreases at a constant rate r and W_t increases at a constant rate i, we have

$$EW = \int_0^\infty P_0 e^{-rt} W_0 e^{it} \, dt = \int_0^\infty P_0 W_0 e^{(i-r)t} \, dt$$

which integrates into $P_0 W_0/(r-i)$ if $r > i$ but integrates into positive infinity if $i > r$.

At present it does not seem reasonable to have an uncertainty discount rate of 1% per annum or higher. (In my view, it should be a very small fraction of 1%, perhaps about 0.01%. A constant discount rate of 1% per annum means that we are less than 37% confident of surviving the next 100 years, and more than 99.99% certain of *not* surviving the next 1,000 years.) On the other hand, world population is growing at more than 1%. Moreover, per capita income is also growing. Though this does not ensure increasing per capita welfare, most happiness surveys indicate so, at least it is not decreasing. Despite the continued presence of famines and wars, their frequency of occurrence has decreased and can be expected to decrease further in the future. Thus it seems that our aggregate welfare is increasing at a rate larger than 1% per annum. Does this mean that our expected welfare is infinite?[3]

The answer is an emphatic "no!". While our population is growing strongly at present and can be expected to do so in the near future, it cannot grow at this high rate indefinitely. Growing at 1% per annum, our population will be increased by 267, 593, 000, 000, 000, 000, 000, 000, 000, 000, 000, 000, 000, 000, 000, 000, 000, 000, 000, 000, 000, 000, 000, 000, 000, 000, 000, 000, 000, *times* just after 20,000 years. At this size, just our own mass would exhaust not only all the mass of the earth and our solar system but also all other masses within a radius of 20,000 light-years. Even if we could travel at the speed of light to colonize other planets, we would not be able to find enough mass just to constitute our bodies. Doubling the time doubles the radius reachable by spaceships travelling at the speed of light. Assuming a more-or-less uniform distribution of mass in space, this increases reachable mass by eight times. On the other hand, growing at a rate of 1% per annum, doubling the time from 10,000 to 20,000 years increases the population by $163{,}582 \times 10^{38}$ times and doubling the time from 20,000 to 40,000 years increases the population by $267{,}593 \times 10^{81}$ times. Thus, such exponential increases must eventually outstrip the mass reachable even at the speed of light.

Using the same argument, we can show that our population cannot increase indefinitely even at 0.001% (or even smaller) per annum, a rate that may be regarded as smaller than a reasonable rate for uncertainty discount. (At a discount rate of 0.001% per annum, we are more than 99% certain that we can survive more than 1,000 years, and more than 90% certain that we can survive more than 10,000 years). It can then be shown that our expected welfare must be finite unless our per capita welfare also increases and that the rate of increase in per capita welfare plus that of population increase add up to be higher than the uncertainty discount.

Can our per capita welfare increase at a rate faster than the rate of uncertainty discount? Ignoring transformation of our own selves through high technology (including advanced forms of genetic engineering), it can be convincingly argued that per capita welfare cannot be increased without limit. Due to biological limitations, the most blissful level of welfare our brain can enjoy is finite. No

[3] I take our instantaneous welfare to be our aggregate welfare (i.e. average welfare per head times the number of individuals) at the relevant time. I have argued strongly for taking this aggregate or total view elsewhere (Ng 1983b, 1986, 1989). However, for the purpose here, if the average view is taken, it really strengthens the argument in the text that our expected welfare is finite.

matter how much and how good objective factors are provided, our subjective appreciation has an upper limit. Thus, unless our subjective self can also be transformed, our per capita welfare must be limited. Hence, our expected welfare must be finite, ignoring the possibility of subjective transformation to the consideration of which we now turn.

IV. The transformation of our own selves through genetic engineering?

Despite the truly spectacular advances in genetic engineering, it is far too early to speak of the transformation of our own selves in operational terms. In fact, if we attempt to do that in the near future, it would probably be disastrous. However, since we are dealing with an issue concerning our welfare in the far future, we cannot ignore the possibility of self-transformation through genetic engineering or perhaps some other means we cannot imagine now, though this can only be discussed in very speculative terms.

Allowing for the possibility of our self-transformation in the future, it cannot be ruled out that our capacity for happiness may be tremendously increased. Perhaps our welfare (per person) could be increased a hundred, a million or even a trillion times. However, such an increase can hardly be infinite. Matter has to be organized in sufficiently complex form of a brain before it can be capable of consciousness. [I understand that this touches on the tricky mind-matter problem called the world knot over which philosophers have pondered over thousands of years without a definite conclusion; but see Ng (1990) for a discussion.] Consciousness is a necessary but probably not a sufficient condition for the perception of welfare and diswelfare. It is not accidental that the brain of *homo sapiens* is bigger than other species. In fact, natural selection should ensure that unnecessary size will not survive the process of competition. Thus we need a minimum size to achieve a given level of brain functioning. Perhaps the efficiency here could be further improved somewhat, but not indefinitely. It might also be possible to increase our capacity for happiness with the given brain size by shifting our brain functions towards the perception of well-being. But the scope for doing this is also limited. Eventually, the ceiling will be hit that increasing the capacity for happiness necessitates a bigger brain. Then our argument above on the physical limitation of the increase in our population size can be used on the limitation on our transformed brain size and hence on our future level of welfare. That our total welfare even in the far future cannot be infinite seems an inevitable conclusion.

If we believe that our expected welfare is infinite, we should be extremely cautious in the sense of rejecting all measures that may lead to however big (but finite) improvements but has some positive (no matter how small) probability of causing our destruction. The argument above that our expected welfare is not infinite releases us from the moral necessity of having to be extremely cautious. An intermediate case applies if we believe that our expected welfare is finite given no self-transformation but is infinite given self-transformation. Then we have to be extremely cautious in all measures except those related to our self-transformation on which we should still be very cautious.

References

Chipman JS (1960) The foundation of utility. Econometrica 28: 193–224
Hawking SW (1988) A brief history of time. Bantam, London
Ng Y-K (1983a) Welfare economics: introduction and development of basic concepts, 2nd Edn., Macmillan, London
Ng Y-K (1983b) Some broader issues of social choice. In: Pattanaik PK, Salles M (eds). Social choice and welfare. North-Holland, Amsterdam, pp. 151–174
Ng Y-K (1984) Expected subjective utility: Is the Neumann-Morgenstern utility the same as the Neoclassicals'? Soc Choice Welfare 1: 177–186
Ng Y-K (1986) On the welfare economics of population control. Popul Develop Rev 12: 247–266
Ng Y-K (1989) What should we do about future generations? The impossibility of Parfit's Theory X. Econ Philos 5: 135–253
Ng Y-K (1990) Utilitarianism and interpersonal comparison: some implications of a materialist solution to the world knot (typescript)
Ng Y-K (forthcoming) Welfarism and utilitarianism: a rehabilitation (Utilitas)