

NUCLEAR DETERRENCE AND FUTURE GENERATIONS

Jefferson McMahan

THE ARGUMENT

People's views about nuclear weapons tend to reflect the ordering of their fears. A crude generalization might be that those whose position is characterized primarily by opposition to nuclear weapons tend to fear nuclear war more than they fear the Soviets, while those who are disposed to support nuclear weapons tend to fear the Soviets more than they fear nuclear war. This in part explains why opponents of nuclear weapons often try to support their position by describing the probable effects of nuclear war and by deemphasizing the Soviet threat, while supporters of nuclear weapons stress the threat from the Soviet Union and the evils of Soviet communism, and maintain that nuclear war would be "survivable." In short, each group seeks to evoke in others the fears that motivate its own position, and to quieten fears that underlie the opposing view. Of course, many defenders of nuclear deterrence may also believe that only the threat of retaliation prevents the Soviets from attacking Western countries with nuclear weapons, and so they could legitimately appeal to the fear of nuclear destruction in support of the retention or acquisition of nuclear weapons. But the fact that they do not ordinarily do this suggests that they recognize that fear of nuclear destruction leads naturally to opposition to nuclear weapons.

The conflict between these two basic attitudes is most strikingly manifest in the old debate about whether it would be better to be "red" or dead. Of course, this debate greatly oversimplifies the issue, since no choice between policies presents us with a stark choice between Soviet domination and nuclear war. But the comparative evaluation of these two outcomes does have a place in one interesting type of

An earlier draft of this essay was read at the California Institute of Technology, the University of Illinois at Chicago, and the University of Chicago. The discussion following each reading was very helpful in my revisions. I have also benefited from written comments by Robin Attfield, Robert McKim, Jan Narveson, and Derek Parfit. My greatest debt is to Steven Lee, whose penetrating comments caused me to make a number of important changes in the argument.

consequentialist argument in which probabilities are also taken into account. For example, someone who believes that it would be better to be red than dead might argue, on the basis of that assumption, that it would be preferable to adopt a policy which, in comparison with the present form of nuclear deterrence, would decrease the probability of nuclear war, even if it would increase the probability of Soviet domination, as long as the increase in the latter probability would not be significantly greater than the decrease in the former.

The assumption about whether it would be better to be red or dead as it would appear in such an argument might be interpreted in different ways, some more defensible than others. The assumption might simply express a personal preference; that is, it might express nothing more than that one would, or would not, find life under Soviet domination worth living. So interpreted, the assumption would be absurdly narrow as the basis for a consequentialist argument. The assumption might, however, be more generously interpreted as expressing a view about whether it would be better for one's country *as a whole* to suffer nuclear destruction or to fall under Soviet domination. But even on this broader interpretation, an assumption of this sort would still be too narrow to serve as a premise in a plausible consequentialist argument, because it would still exclude from consideration the interests of persons in other countries—both enemy countries and countries not directly involved in the conflict. The interests of Americans are not the only interests that count. Nor, indeed, are existing generations the only generations that count. Thus, perhaps the most important reason why an assumption about whether it would be better to be red or dead is too narrow is that it fails to take account of the effects of nuclear war on future generations.

For these reasons, a consequentialist argument against nuclear deterrence based solely on the claim that it would be better to be red than dead would be implausible. But an analogous argument can be constructed on a broader foundation. The foundation would still consist in a comparison between the two outcomes: nuclear war and Soviet domination. But the comparison would take into consideration the interests of existing people everywhere, and would also take account of how future generations might fare in each of the two outcomes. The main aim of this essay is to present and discuss this analogous argument and, in particular, to defend its central assumption: that it is of the utmost moral importance to ensure the existence of future generations.

The argument, stated fully, is as follows:

- 1a. Nuclear war would be a greater evil than Soviet domination where future generations are concerned.
- 1b. Nuclear war would be worse than Soviet domination for people in the US and for people in allied countries which would also be directly involved in the war.
- 1c. Nuclear war would be worse than Soviet domination for people in countries not directly involved in the war, and it would also be worse for the Soviets and their allies.
2. It is therefore considerably more important to prevent nuclear war than it is to prevent Soviet domination.
3. The present policy of nuclear deterrence has a high probability of preventing Soviet domination, but it also has a significant probability of leading to nuclear war, even with the relatively near future—say, within the next 30 years.

4. The abandonment of nuclear deterrence by the US and the adoption of a policy of nonnuclear defense would decrease the probability of nuclear war, but would also increase the probability of Soviet domination. It would not, however, increase the latter probability by significantly more than it would decrease the former. Indeed, it might increase the latter by *less* than it would decrease the former.

5. Any other negative effects of this change in policy, even when considered together with the increase in the probability of Soviet domination, would not be sufficient to outweigh in importance the decrease in the probability of nuclear war.

6. One of two policies is superior if it would, when compared with the other policy, reduce the probability of a very bad outcome, even if it would also increase the probability of a significantly less bad outcome, provided that the increase in the latter probability would not be significantly greater than the decrease in the former, and provided that the adoption of the policy would not have other undesirable effects which, together with the increase in the probability of the less bad outcome, would outweigh the advantage of decreasing the probability of the very bad outcome.

7. A policy of unilateral nuclear disarmament by the US, combined with a policy of nonnuclear defense, would be superior to continued reliance on the present policy of nuclear deterrence.

(This argument also applies, *mutatis mutandis*, to the Soviet Union; but, since it is intended as a contribution to an essentially Western debate, it is preferable to direct the argument against the US, whose policies we are in a better position to influence.)

It is important to notice that the conclusion of the argument is *not* that unilateral nuclear disarmament would be the best policy for the US to adopt. That may well be true, but it is not entailed by the argument. The conclusion of the argument is only that unilateral nuclear disarmament is preferable to the present policy of nuclear deterrence. This leaves open the possibility that there may be some further alternative that would be superior to both—for example, a policy of minimal deterrence (that is, a policy that would aim to deter a nuclear attack, and nothing else, with the smallest arsenal that would be sufficient to guarantee deterrence).

The argument rests on a number of controversial claims, each of which might be challenged. My own view, however, is that each is defensible. Although I shall argue at some length in support of premise 1a, spatial constraints make it impossible to provide a thorough defense of any of the premises. I hope to say enough in support of each premise to show that the argument as a whole is indeed plausible. But my main aim must be simply to illuminate the structure of the argument, and to suggest what considerations are relevant to the assessment of the premises.

CONSEQUENCES FOR PRESENT GENERATIONS

It is primarily premise 1a—that nuclear war would be worse than Soviet domination where future generations are concerned—that makes this argument against nuclear deterrence distinctive. As we shall see, this initial premise is based on the assumption that it is considerably more important to ensure that future generations exist than to ensure that, if they exist, they will not exist under Soviet domination. Since my argument in support of this assumption is lengthy, controversial, and more distinctively philosophical in character than the discussions of the other premises of

the argument, it merits a separate section to itself, which I shall reserve until after I have discussed and defended the other premises of the main argument; for it will help to see how the argument as a whole works before we focus on the justification of this one premise.

It might be suggested that premises 1b and 1c, which together assert that nuclear war would be worse for existing people than Soviet domination, would themselves be sufficient to establish premise 2—the claim that it is considerably more important to prevent nuclear war than to prevent Soviet domination. In that case, it might be thought that premise 1a would be superfluous. That would be a mistake. Premises 1a, 1b, and 1c work together to support premise 2. The more support there is for each of the first three premises, the stronger premise 2 will be—that is, the greater will be the relative importance of preventing nuclear war over preventing Soviet domination. And the greater the relative importance of preventing nuclear war is, the more readily and decisively will the argument carry through to its conclusion. (For the greater the relative importance of preventing nuclear war, the greater will be the increase in the probability of Soviet domination that we would accept in exchange for a fixed decrease in the probability of nuclear war.)

Bearing in mind this fact about the structure of the argument, let us provisionally grant premise 1a, and go on to consider premise 1b—the claim that nuclear war would be worse than Soviet domination for people in the US and allied countries. This claim is supported by several lines of argument. Since any two-sided nuclear war would be unlikely to remain limited,¹ nuclear war would probably result in the deaths of most people in the US and in allied countries. If most people would find life under Soviet domination worth living, then it follows that, for most of those who would be killed in a nuclear war, nuclear war would be worse than Soviet domination. And indeed it seems plausible to suppose that most people in the West would find life under Soviet domination worth living, just as most people do who presently live under Soviet domination.

Suppose that, contrary to this conjecture, many people would personally prefer death to life under Soviet domination—that is, they would not find life under Soviet domination worth living. Would this show that Soviet domination would be worse than nuclear destruction? Not necessarily; for the view that Soviet domination would be worse than nuclear destruction would presumably be based on a claim about the overriding importance of liberty, but the view is itself less attentive to the demands of liberty than the contrary view. It ignores the preferences of those who would prefer life under Soviet domination to death. But the view that nuclear destruction would be worse respects the preferences of those who would prefer death to life under Soviet domination, for Soviet domination would allow these individuals to choose death rather than submit to domination by the Soviets.² In short, while nuclear war would be worse for those who would prefer life under Soviet domination to death, Soviet domination need not be worse for those who would prefer death to life under Soviet domination.

What about those persons in the US and allied countries who might survive a nuclear war? Would life among the radioactive ruins of the US and Western Europe be preferable to life under Soviet domination? What primarily distinguishes life in many Western countries today from life in Soviet-dominated countries is that people in the West enjoy political, social, and economic freedoms denied to people in the East. These would not survive a nuclear war. Michael Howard has claimed that, from the point of view of the survivors of a nuclear war, “the political, cultural and

ideological distinctions that separate the West from the Soviet Union today would be seen, in comparison with the literally inconceivable contrasts between *any* pre-atomic and *any* post-atomic society, as almost insignificant."³ His main reason for this claim is that any regime that might emerge in the US after a nuclear war would be inescapably authoritarian, and thus would, in this crucial respect, resemble the regime that governs the Soviet Union today.

What about those who, in spite of this, would prefer life amid the postwar ruins to life under Soviet domination? The short answer is simply that, while Soviet domination would of course deny them their preference, it is by no means clear that nuclear war would satisfy it, for survival in nuclear war cannot be guaranteed.

Together these arguments provide strong support for the claim that nuclear war would be worse than Soviet domination for people in the US and in allied countries. Let us turn, then, to the claim that nuclear war would also be worse for people elsewhere in the world. It is, of course, completely uncontroversial that nuclear war would be worse than Soviet domination for people in the Soviet Union and allied countries. It is, however, less clear that nuclear war would be worse than Soviet domination for people in countries not directly involved in the war. One reason why this comparison is more difficult is that it is uncertain what the effects of nuclear war would be in neutral countries. It is arguable that nuclear war would result in the deaths of most, if not all, of the people in these countries.⁴ In that case nuclear war would be worse for these people than Soviet domination, for the same reasons that it would be worse for people in the US and allied countries. But suppose that the effects would be more limited—for example, that they would consist primarily in a certain number of immediate deaths from fallout, and in environmental damage, increased rates of cancer and birth defects, and severe political and economic dislocation. Even in this case the upheaval in these countries would be enormous, and would be likely to give rise to authoritarian political structures no less repressive than those in Soviet-dominated countries today. Thus it is difficult to believe that life under these conditions would be better than life under Soviet domination.

For many of the world's people, life in a world dominated by the Soviet Union might be no worse than life is at present. It is difficult to substantiate this claim, however, partly because it is difficult to predict what a world dominated by the Soviet Union would be like. Conditions would undoubtedly be different in different places. But, if it is permissible to extrapolate from our knowledge of life today in Soviet-dominated countries, the claim has some plausibility. Life in Soviet bloc countries today seems, in general, no worse than life in many other countries—some of which are effectively under American domination. To see the force of this point, one might ask oneself, for example, whether one would prefer to be a dissident in Poland today or to have been a dissident in Iran under the Shah. While American governments have shown far greater respect than Soviet leaders for the rights and liberties of their own citizens, their concern for the rights and liberties of citizens of Third World client states has been no greater than that shown by Soviet leaders toward the citizens of Soviet client states.⁵ To say this is not to present an apology for Soviet efforts to control the affairs of other countries, but only to provide a *comparative* evaluation of the possible consequences of Soviet domination for people outside the US and Western Europe.

The foregoing arguments provide strong support for the claim that nuclear war would be worse for existing people than Soviet domination. If we provisionally grant that nuclear war would also be worse where future generations are concerned,

then it seems clear that nuclear war would be a far greater evil than Soviet domination, and therefore that it is considerably more important to prevent nuclear war than to prevent Soviet domination. In other words, premise 2 is true.

Let us turn, then, to premise 3. Defenders of the present policy of nuclear deterrence will not, of course, object to the claim that this policy has a high probability of preventing Soviet domination. But they will deny that there is a significant probability of nuclear war in the relatively near future under nuclear deterrence as it is currently practiced. Strictly speaking, the argument does not actually require the claim that nuclear deterrence is likely to lead to nuclear war; all it requires is the claim that an alternative policy would have a *lower* probability of leading to nuclear war, without having a significantly higher probability of leading to Soviet domination. Nevertheless, there are several reasons for including the claim that nuclear deterrence as it is presently practiced is likely to lead to nuclear war. One is that, if correct, the claim provides the argument with additional urgency. Premise 6, in particular, seems most cogent in cases in which the worse outcome is highly probable under the alternative policy. Another reason for including the claim is that the arguments that support it will help to show why the shift to a policy of nonnuclear defense would greatly reduce the probability of nuclear war. A brief critique of the present policy of nuclear deterrence will set the stage for a defense of an alternative policy. But the most important reason for including the claim that the present policy of nuclear deterrence has a significant probability of leading the nuclear war is that, without this claim, the argument would be less likely to go through. For, if the probability of nuclear war under nuclear deterrence were very low, then there would be little scope for reducing it by shifting to a policy of nonnuclear defense. And since, as I have conceded, the probability of Soviet domination under nuclear deterrence is low, the scope for increasing it by shifting to a policy of nonnuclear defense is considerable. Given this combination of background conditions, it might seem more reasonable to suppose that the shift to a policy of nonnuclear defense would increase the probability of Soviet domination by significantly more than it would reduce the probability of nuclear war.

Nuclear deterrence as it is presently practiced has a significant probability of leading to nuclear war, for at least five reasons. All five are familiar, and thus require little elaboration here.

1. The practice of deterrence has led (some say necessarily) to the development by both sides of counterforce weapons and the adoption of counterforce strategies. The characteristics necessary for counterforce weapons are the same as those necessary and sufficient for first-strike weapons. Thus, each side's deployments now pose a threat to the other's retaliatory capability. As each develops its counterforce capability, it will have an increasing incentive to use this capability in a first strike of its own. The fears generated by each side's sense of ever-increasing vulnerability, combined with other pressures to insure that deterrence remains "credible," in turn lead to yet more destabilizing new deployments.

2. The practice of nuclear deterrence leads inexorably to competition in the development of new technologies. Each side feels it must press forward with its own research; otherwise the other side might arrive unilaterally at technological discoveries that would provide a decisive strategic advantage. This technological side of the arms race leads not only to fears and suspicions about the sinister projects the other side may be pursuing, but also to the development of dangerously

destabilizing technologies—both of which aggravate the threat of war. For example, early research on ballistic missile defenses led to the development of MIRV's, which have in turn (in conjunction with increases in missile accuracy) given rise to fears of a first strike.

3. The possession of nuclear weapons in large numbers makes it not unlikely that a nuclear war will start by accident or mistake. This could happen in various ways. The most likely is that one side might launch its missiles under the mistaken impression that it is itself under attack. False alarms are in fact a disturbingly frequent occurrence.⁶

4. Nuclear deterrence both sanctions and encourages nuclear proliferation, and nuclear proliferation increases the probability of nuclear war. Nuclear deterrence sanctions proliferation because most of the arguments supporting the possession of a nuclear arsenal for purposes of deterrence are universal in their application; they do not refer to unique features of any particular country's situation. Deterrence also encourages proliferation in at least two ways. The possession of nuclear weapons for purposes of deterrence provides continuing testimony to their value. It also poses a threat to the security of nonnuclear countries, which these countries, following the logic and the example of the threatener, may attempt to meet by developing nuclear arsenals of their own. Thus, fear of the US and the Soviet Union prompted China to develop its arsenal, and fear of the Chinese arsenal at least partly motivated the development of India's nuclear capability. Fear of the Indian capability has now spurred Pakistan's efforts to acquire a nuclear arsenal.

5. Nuclear deterrence stimulates and, indeed, requires the mutual fear and hatred that are more likely than anything else to lead eventually to war. Politicians often find it necessary, as Senator Vandenberg once advised President Truman, "to scare the hell out of the country" with visions of the "communist menace" to whip up public support for expanded arms programs. Even in the absence of this calculated manipulation of our fears, we would naturally fear and hate those who perpetually threaten to incinerate us. But we may also find it psychologically necessary to hate anyone whom we threaten with annihilation. Our consciences may require that we cast our potential victims in evil and dehumanizing roles, as President Reagan did when he described the Soviet Union as "an evil empire" and claimed that communism is "the focus of evil in the modern world."⁷ Under deterrence, these attitudes cannot be significantly relaxed. Thus, E.P. Thompson writes that, "by maintaining each part in a posture of menace to the other, nuclear deterrence fixes indefinitely the tension which makes the resolution of differences improbable."⁸

With these five points as background, let us now turn to the defense of premise 4. This premise makes several factual claims. The first is that the adoption of a policy of nonnuclear defense would reduce the probability of nuclear war. A nonnuclear defense policy would require complete unilateral nuclear disarmament and a reliance on various types of nonnuclear forces for the country's defense. It might, for example, involve greatly strengthening the US's conventional forces; it might involve the reintroduction of conscription (with, one hopes, more generous provisions for conscientious objection than in the past); it might involve the formation of a territorial militia; or it might involve a greatly expanded civil defense program. These are just some of the possibilities. The important point is that, while such a policy would require unilateral nuclear disarmament, it would *not* leave the US defenseless.

The adoption of a nonnuclear defense policy would reduce the probability of nuclear war, for at least six reasons. Five of these refer back to the reasons why nuclear deterrence as presently practiced is likely to lead to nuclear war.

1. The adoption of a nonnuclear defense policy would reduce the probability of a preemptive first strike virtually to zero. Nuclear disarmament would deprive the U.S. of the *ability* to launch a first strike, and it would deprive the Soviet Union of its primary *target* for one. It is sometimes suggested that the Soviets might be tempted to attack the US to prevent the possibility of American nuclear rearmament. But to achieve this aim they would have to attack and destroy all the many nuclear reactors in the US, and they would surely be deterred from such an attack by the fact that it would produce a vast amount of global fallout. It might also be objected that American nuclear disarmament would remove the major obstacle to a Soviet preemptive strike against the Chinese. But, if the Soviets were tempted to attack the Chinese, they would presumably be restrained more by an awareness of the damage such an attack would do to their international reputation than by fears of American retaliation, which would be extremely unlikely.⁹

2. A policy of nonnuclear defense would slow the competition in the development of nuclear technologies by removing one of the two main competitors.

3. A nonnuclear defense policy would greatly reduce the probability of an accidental nuclear war. It would eliminate the possibility of an accidental or unauthorized firing by the US, and would also eliminate the possibility that an accidental firing by the Soviets would lead to retaliation and uncontrollable escalation.

4. A policy of nonnuclear defense might have an inhibiting effect on proliferation. It would presumably restrict the transfer of nuclear materials from the US to other countries. Since it would also fulfill the US's obligations under the Non-Proliferation Treaty, it might help to revitalize the treaty. Finally, just as the present policy of nuclear deterrence testifies to the value of nuclear weapons, so the abandonment of nuclear weapons by the US would testify to the dangers that attend the possession of nuclear weapons.

5. The adoption of a policy of nonnuclear defense would serve to reassure the Soviets about American intentions, and so would help to dispel the fears, suspicions, tensions, and animosities that presently constitute the single most important factor dragging the world toward war.

6. The adoption of a policy of nonnuclear defense could lead the Soviets to reciprocate by greatly reducing the number of their nuclear weapons. This might occur as the natural result of the relaxation of tensions consequent upon the change in American policy. Or, more cynically, it might happen because the Soviets would not wish to be seen as relentlessly militaristic in comparison with the US. They would, in any case, have little to lose by dismantling a large number of their weapons for, given that a significant percentage of their nuclear weapons are at present targeted on American missile silos and bomber bases, the elimination of the US's nuclear arsenal would deprive most of the Soviet missiles of their targets.

At this point it may be objected that these arguments either ignore the deterrent effect of the possession of nuclear weapons, or else assume that Soviet intentions are wholly benign. A policy of nonnuclear defense would deprive the US of the ability to deter a nuclear attack, and thus, it might be argued, would *increase* the probability of nuclear war rather than decrease it.

In response to this objection we need to ask what reason the Soviets would have for a nuclear attack on the US if the US had renounced the possession of nuclear weapons. Since the US would no longer pose an offensive threat to the Soviet Union, they could have no defensive reason for attacking, and the idea that the Soviets might attack out of sheer malice seems excessively cynical. The most plausible suggestion is that they might make selective nuclear strikes in an attempt to achieve domination or conquest. In particular, actual or threatened nuclear strikes might be used in an effort to coerce the surrender of the US, and thereby to establish the global dominance of the Soviet Union.

It would seem, then, that what is really being claimed when it is said that the abandonment of nuclear deterrence would increase the probability of nuclear war is that the abandonment of nuclear deterrence would increase the probability of Soviet nuclear blackmail. But the limited, coercive nuclear strikes the Soviets *might* make in an effort to achieve domination over the US would not constitute nuclear war in the sense intended in the premises of the main argument. These premises presuppose a conception of nuclear war such that nuclear war would involve the relatively extensive use of nuclear weapons, and would result in widespread destruction. Otherwise—if one or two nuclear strikes would count as a nuclear war—the first four premises of the argument would presumably be false.

Given this understanding of what counts as nuclear war, the adoption of a policy of nonnuclear defense would clearly reduce the probability of nuclear war: it would greatly reduce the probability of the extensive use of nuclear weapons. The real objection to the adoption of a nonnuclear defense policy, then, is that it would increase the risk of nuclear blackmail, and thus would increase the probability of Soviet domination. Since the US would certainly be able to defend its borders under a policy of nonnuclear defense, Soviet domination could probably be achieved only through some form of nuclear coercion, involving either actual or threatened nuclear destruction. Thus, a limited amount of nuclear destruction in the US should be counted among the possible costs of Soviet domination. (It might be objected that, if we assume that Soviet domination might involve a limited amount of nuclear destruction, this will weaken the case for premises 1a through 2. I think, however, that careful reflection on the arguments for those premises will show that they are not undermined by this assumption. If Soviet domination would involve a limited amount of nuclear destruction, that would of course make it worse than it would otherwise be, precisely because it would then involve some of the evils characteristic of an even worse outcome—all-out nuclear war.)

How likely would it be that the US, if it were to give up nuclear weapons, would be subjected to nuclear blackmail leading, if successful, to Soviet domination? The answer depends in part on our assessment of the motives and goals of the dominant Soviet leaders. If, as is arguable, Soviet military policies are motivated primarily by defensive concerns, then the abandonment of nuclear deterrence would not significantly increase the probability of nuclear blackmail or Soviet domination. It would, however, be unwise to rest the argument on sanguine assumptions about Soviet motivations. So we should assume that the familiar allegations about the Soviets' aggressive designs have some plausibility. And we must therefore concede that the abandonment of nuclear deterrence might significantly increase the probability of nuclear blackmail and Soviet domination. But, even if the Russians would no longer be restrained by fear of nuclear retaliation, they would still have other reasons for refraining from attempting nuclear blackmail. One is that nuclear

threats or an actual nuclear attack would almost certainly provoke an effort at nuclear rearmament by the US. Moreover, the Soviets could not hope to subdue the entire world simultaneously, and other countries with a potential nuclear capability might be frightened into developing their own nuclear arsenals, which would then be arrayed against the Soviet Union. For an attempt at subjugating the US through nuclear blackmail, the Soviet Union would probably pay a high price in the creation of a large number of determined enemies. Finally, even if the effort at nuclear blackmail were to succeed, the Soviets would still face what I think would be predictably insurmountable problems in trying to subjugate and control the United States, a large country, both geographically and in terms of population, whose population is remarkably united in its hostility to Soviet communism.¹⁰

In short, while shifting to a policy of nonnuclear defense would undoubtedly increase the probability of nuclear blackmail and Soviet domination, how much it would do so is a matter of speculation. But, given that the shift to a policy of nonnuclear defense would *greatly* reduce the probability of large-scale nuclear war, it is hard to believe that it would increase the probability of Soviet domination by *significantly more* than it would decrease the probability of nuclear war. If it would not, then that is all that the argument requires.

We have so far compared the present policy of nuclear deterrence with a policy of nonnuclear defense with respect to two possible outcomes: nuclear war and Soviet domination. Either or both of these policies, however, might lead to other possible outcomes. If we were to compare the two policies with respect to all the possible outcomes, the argument would become hopelessly complex. Thus premise 5 has been introduced to eliminate the need to conduct a series of detailed comparisons. It asserts that the apparent superiority of a policy of nonnuclear defense over the present policy of nuclear deterrence where nuclear war and Soviet domination are concerned is not outweighed by the superiority of nuclear deterrence where other possible outcomes are concerned.

The strongest challenge to this premise lies in the claim that the shift to a policy of nonnuclear defense would greatly increase the probability of conventional war. While it is unlikely that the Soviet Union would be tempted to launch a conventional attack against the US, it might be less inhibited about initiating conventional war elsewhere in the world—for example, in Western Europe. At present the possibility of escalation to all-out nuclear war helps to deter any type of attack on the US or its allies. Since the adoption of a policy of nonnuclear defense would eliminate this threat, the risks of conventional aggression might then appear acceptable.

This is an important challenge, but I think it can be met. There are three relevant points. One is that the possibility of American nuclear rearmament would have a deterrent effect. But a more important point is this. There are two ways to dissuade a potential aggressor from attacking. One is to threaten him with punishment if he does attack, so that the benefits he might derive from attacking would be outweighed by the harm he would suffer. Nuclear deterrence attempts to influence a potential attacker's calculations in this way. The other way is to arrange things so that, whatever his aim in attacking, he would be unable to achieve that aim. This second form of dissuasion is often referred to as "deterrence by gain denial" or "defensive deterrence."

Unlike strategic nuclear attacks, conventional attacks can be effectively deterred by defensive means. Indeed, the threat of "gain denial" is in general a more *credible* deterrent to conventional attack by a nuclear-armed aggressor than the threat of

retaliation, since defensive measures do not involve the same risks that retaliation involves. Already there is a strong movement of opinion in Western Europe favoring a policy of defensive deterrence in Europe. Since the total gross domestic product of the European NATO countries alone is substantially greater than that of the Warsaw Pact (including the Soviet Union), there is no doubt that collaboration between the European NATO countries and the US could lead to the development of nonnuclear forces amply sufficient to defeat, and therefore to deter, a conventional Warsaw Pact invasion.

At this point it may be objected that the Warsaw Pact could overcome NATO's nonnuclear defenses through the use of tactical nuclear weapons, so that the likeliest consequence of America's abandonment of nuclear deterrence would not be a conventional attack on Europe, but a *combined* conventional and tactical nuclear attack. This objection assumes that tactical nuclear weapons provide strong or perhaps even decisive tactical advantages. This assumption seems to me to be false. Spatial constraints prevent me from challenging it here, although I have done so elsewhere.¹¹

The third point that can be made in response to the charge that American nuclear disarmament would increase the probability of conventional war is in part an *ad hominem* point. Many of those who argue that nuclear weapons are needed to deter conventional war claim that the deterrence of conventional war is nearly as important as the deterrence of nuclear war, for modern conventional war would be only marginally less terrible than nuclear war. But, if it is true that conventional war would be that destructive, then it follows that the threat of conventional destruction should be virtually as effective a deterrent to conventional war as the threat of escalation to nuclear war.

The possibility of American nuclear rearmament, the threat of effective "gain denial," and the terrible destruction potential of modern conventional warfare could together provide a strong deterrent to conventional war. Thus, if the adoption of a policy of nonnuclear defense would increase the probability of conventional war, it need not do so by much.

Even if we assume that the adoption of a policy of nonnuclear defense would have other negative effects than the increase in the probability of Soviet domination, these other effects would at least in part be cancelled out by certain other positive effects of the change in policy. Among these positive effects would be a significant decrease in the production of nuclear wastes, a significant reduction of the probability of a serious nuclear accident, and the elimination of the threat posed to civil liberties and democratic institutions by the possession of nuclear weapons. All things considered, it seems likely that a policy of nonnuclear defense could satisfy the condition stated in premise 5.

As I have suggested, premise 5 is intended to deal with the fact that negative outcomes other than nuclear war and Soviet domination are relevant to the comparison between nuclear deterrence and nonnuclear defense. At this point it may be objected that my argument gives insufficient attention to the main *positive* outcome relevant to the comparison of the two policies—the continuation of the status quo, or the avoidance of *both* nuclear war and Soviet domination. For it is surely relevant to the comparison of the two policies to ask which has the greater probability of avoiding both these disastrous outcomes.

In an important paper on the subject of nuclear deterrence, Gregory Kavka proposes and defends a principle he calls the Disaster Avoidance Principle. The principle asserts that

when choosing between potential disasters under two-dimensional uncertainty [that is, "the chooser has no reliable quantitative estimates of the relevant utilities and probabilities, but has confidence in his judgment of their ordinal rankings"], it is rational to select the alternative that minimizes the probability of disaster occurrence.¹²

Kavka claims that this principle supports nuclear deterrence over unilateral nuclear disarmament (or nonnuclear defense) since, of the two policies, nuclear deterrence offers the greater probability of avoiding both relevant disasters, nuclear war and Soviet domination. He claims that, because of this, the policy of nuclear deterrence is superior, even though it has a greater probability of leading to the worse of the two disasters. This argument directly challenges the plausibility of the principle stated as premise 6 in my argument.

Kavka's claim that nuclear deterrence offers the greater probability of maintaining the status quo is based on the assumption that the probability of nuclear war under nuclear deterrence is less than the probability of Soviet domination would be under unilateral nuclear disarmament. I think that my earlier arguments show that this assumption is very much open to doubt, although it is compatible with the main argument of this essay. But, even supposing that his assumption is correct, it is not sufficient to establish that nuclear deterrence offers the greater probability of disaster avoidance. To establish this conclusion, it is necessary to make the further assumption that the probability of nuclear war under unilateral nuclear disarmament would be greater than (or at least not less than) the probability of Soviet domination under nuclear deterrence. For, if the probability of nuclear war under unilateral nuclear disarmament would be *less* than the probability of Soviet domination under nuclear deterrence, then, given Kavka's assumption that there are no reliable quantitative estimates of the relevant probabilities, we could not know which of the two policies would provide the greater probability of disaster avoidance.

It is not implausible, moreover, to suppose that the probability of nuclear war under unilateral nuclear disarmament *would* be less than the probability of Soviet domination under nuclear deterrence. I have already suggested why I think that the probability of nuclear war under unilateral nuclear disarmament would be very low. And, although I have conceded that the probability of Soviet domination under nuclear deterrence is low, even the most ardent proponents of nuclear deterrence have argued that the probability is not insignificant. They have argued that any number of apparently insignificant weaknesses in the American strategic position could give the Soviets a psychological advantage and enable them to coerce the US to surrender.¹³

Of course, nothing I have said so far has been sufficient to prove the falsity of either of the assumptions necessary to show that nuclear deterrence offers a greater probability of disaster avoidance. But I think enough has been said to show that neither assumption is obviously true. This being the case, it is not clear whether Kavka's Disaster Avoidance Principle is actually applicable to the comparison between nuclear deterrence and unilateral nuclear disarmament (or whether, if it is applicable, it actually favors nuclear deterrence). This suggests that, in comparing these two policies, it is *at least* equally plausible to appeal to my premise 6 as it is to appeal to the Disaster Avoidance Principle.¹⁴

A further objection to my argument that might be mentioned here is that, since each of us would (one hopes) be willing to risk death to defend his country,¹⁵ we

should all be willing to risk nuclear destruction by supporting nuclear deterrence, in order to prevent Soviet domination. Yet, there are important differences between risking individual death and risking collective death. One is that to risk collective death is to put future generations at risk as well. If it matters whether future generations will exist, then each individual may be willing to risk his own life only on the assumption that failure will not entail the death of all. Thus from a willingness of each to risk his own life we cannot infer a willingness of all to risk the lives of all.

The foregoing arguments, while not conclusive, should be sufficient to show that premises 1b through 6 are not obviously implausible. It remains to be shown that premise 1a is also defensible.

CONSEQUENCES FOR FUTURE GENERATIONS

The main reason for thinking that nuclear war would be worse than Soviet domination where future generations are concerned is that nuclear war could lead to the extinction of the human race, and it is considerably more important to ensure that future generations will exist than to ensure that, if they exist, they will not exist under Soviet domination. Of course, it is by no means certain that nuclear war would lead to extinction. On the contrary, most scientists seem to agree that extinction would be unlikely. On the other hand, it seems uncontroversial that a large-scale nuclear war (which is probably the only possible kind of nuclear war) would greatly reduce both the number of future people and the future quality of life. It will become evident that, because of this, the explanation I shall offer of why extinction would be a terrible tragedy will also imply the truth of premise 1a even on the assumption that nuclear war would not lead to extinction. Thus my argument in this section will support the claim that nuclear war would be worse than Soviet domination where future generations are concerned regardless of whether or not nuclear war would lead to extinction—although the extent to which nuclear war would be worse than Soviet domination will be greater the more likely it is that nuclear war would lead to extinction.

The claim on which premise 1a rests, and which I hope to establish in this section, is that it is of the utmost moral importance to ensure the existence of future generations. Many people feel intuitively that this claim is correct, but are nevertheless unable to find arguments to support their conviction.¹⁶ Others believe the only moral reasons for ensuring the existence of future generations are those that arise from the fact that we, the living, prefer that there should be future generations—because, for example, we desire to have children.¹⁷ After all (it might be argued), preventing the existence of future generations could not be worse for future generations themselves, for nothing can be worse for people who never exist; so preventing their existence would not be against *their* interests or violate their rights. If those who argue in this way are right and the only reasons we have for ensuring the existence of future generations are ones that appeal to our own interests, then premise 1a will contribute nothing to support the claim that nuclear war would be worse than Soviet domination which is not already provided by the claim that it would be worse for existing people.

This challenge to premise 1a might be developed in the following way. It is only on the condition that future generations will exist that we can affect their interests. Thus, a concern for their interests cannot provide a reason for ensuring their existence. On the other hand, a concern for their interests can and, indeed, does provide

a reason for ensuring that, *if* they exist, they will not exist under Soviet domination. In the words of one recent writer, "if we believe that the risks of deterrence are worth taking for ourselves than we need not shrink from taking them . . . for our descendants."¹⁸ This is so (it might be argued) because, while *we* have much to lose if the gamble fails and the race is exterminated, future generations would, strictly speaking, lose nothing at all.

This argument rests on a mistake that has been exposed by Derek Parfit.¹⁹ Suppose that, as a result of American nuclear disarmament, the Soviets were to dominate the world. This would obviously have widespread effects on people's lives everywhere. There would be important contrasts between life under Soviet domination and life as it would otherwise have been. One important contrast, often ignored, is that different people would meet, and different marriages would be made, so that different children would be born. Even in those marriages that would be the same, children would be conceived at different times and thus would develop from different genetic materials. As Parfit has shown, this would in fact be sufficient to make them different children. In short, if the Soviets came to dominate the world, this would dramatically affect who would subsequently exist. As time passed, the proportion of people who would not have existed had the Soviets not dominated the world would increase until eventually there would be no one in existence who would also have existed had the Soviets not dominated the world. There would, therefore, be relatively few future people of whom it could be said that they were affected for the worse by American nuclear disarmament and the subsequent domination of the world by the Soviets. Assuming that the others would have lives worth living, they could not claim to have been affected for the worse by Soviet domination, since, were it not for Soviet domination, *they* would not have existed.

This shows that, in the case of most future people, our reason for ensuring that they will not exist under Soviet domination cannot derive from a concern for their interests. Our reason must instead be more impersonal in character. But from a more *impersonal* point of view it must surely be more important that future generations exist, if their lives would be worth living, than that they do not exist under Soviet domination—especially since Soviet domination could not be expected to last forever.

I shall reinforce this conclusion with several arguments for the claim that, while preventing the existence of future generations would not be against their interests, it is nevertheless of the utmost moral importance not to prevent their existence. One such argument appeals to the fact that our lives would be impoverished by the expectation that we will be the final generation. At present our lives are enriched by the assumption that they will be linked in various ways with the lives of future people. We rely on future generations for the furtherance and completion of projects we have begun or taken over from our ancestors; we depend on them to preserve and enrich our culture, and to help fulfill our ideals; and we hope that they will benefit from and appreciate our works, providing us with posthumous recognition. If we were to suppose that there would be no future generations, many of our present activities would be robbed of much of their meaning.²⁰

These are undoubtedly important reasons for ensuring the existence of future generations. Again, however, if the force of these points is *only* that it would be worse for existing people if there were to be no future generations, then these points will contribute nothing to the larger argument against nuclear deterrence that is not already provided by premises 1b and 1c. It is, however, equally plausible to sup-

pose that there is *independent* value in, say, the evolution of our culture, so that it is important for our culture to continue to develop quite apart from the fact that *our* lives would be impoverished by the belief that the evolution of our culture were at an end. If this further claim is accepted, we have a reason for ensuring the existence for future generations that is independent of the interests of existing people.

Another and perhaps stronger argument for the claim that it is morally important to ensure the existence of future generations also makes no appeal to the interests of existing people. This argument moves from the claim that there is a principle of non-maleficence that provides a moral reason not to bring a person into existence if his life would be worse than no life at all, or "worth not living," to the claim that there is a principle of beneficence that provides a moral reason to bring a person into existence if his life would, on balance, be worth living. The argument takes as its first premise the claim that it would be wrong, other things being equal, to bring a person into existence if his life would predictably be worth not living. This seems uncontroversial. But how can we best explain *why* it would be wrong? It is tempting to appeal to side-effects, to the fact that it is normally worse for existing people if a person who is utterly wretched comes to exist. But this explanation is excluded by the *ceteris paribus* clause. And in any case the appeal to side-effects could provide only a partial explanation of why it would be wrong to bring a miserable person into existence. For it is only contingently true that it is worse for existing people when miserable people come into existence. There could be cases in which this would be better for existing people.

A second possible explanation is that to bring a miserable person into existence is worse *impersonally*, since it would involve a net increase in the amount of misery in the world. The argument could in fact be run on the basis of this explanation, and the conclusion would be substantially the same (although it would differ in form). But this second explanation seems less plausible than a third: that it is wrong to bring a person into existence if his life would be worth not living simply because to do so would be bad for that person—not just impersonally bad, or bad "from the point of view of the universe," but bad from the point of view of the person himself. To bring such a person into existence would be to *harm* that person. Of course, in order to defend the view that it would be wrong to bring a person into existence if his life would be worth not living, we need more than the simple claim that to bring such a person into existence would be to harm him. We also need the further claim that, other things being equal, it is wrong to do what will harm people. This too seems uncontroversial. Most people accept as part of their morality a principle of non-maleficence.

The next stage in the argument is to point out that, if to bring into existence a person whose life is worth not living is to harm that person, then to bring into existence a person whose life *is* worth living must be to benefit that person.²¹ If, in addition to a principle of non-maleficence, we also accept a principle of beneficence, then it follows that there is a moral reason to benefit people by bringing them into existence.

Now, there are several ways of resisting this conclusion which are compatible with the assumption that to bring into existence a person whose life is worth living is to benefit that person. The most obvious is to deny that there is a moral reason to benefit people. If there is a moral reason not to harm people but no moral reason to benefit people, then it follows, given our assumption that to be brought into existence can be either a benefit or a harm, that it would be wrong, other things being

equal, to bring a miserable person into existence, but not wrong not to bring into existence a person whose life would be worth living.

While it is commonly assumed that it is morally more important to prevent or alleviate misery than to promote happiness, the claim that there is *no* moral reason to promote happiness—*no* general reason to benefit people—seems unacceptably strong. Among other things, it implies that it would not be wrong to fail to prevent a dramatic decline in the quality of life, as long as the decline involved only the loss of certain sources of happiness, and not an increase in suffering or misery. It implies that this would not be wrong even if the decline could be prevented at little or no cost to the agent. This seems unacceptable.²²

Another way of denying that the expectation that a person would have a life worth living provides a moral reason for bringing him into existence is to appeal to the view that there is a moral asymmetry between *doing* and *not doing*—the view that we are morally more responsible for what happens as a result of what we do than we are for what happens as a result of what we do not do. If a person's coming into existence can be either good or bad for that person, then the coming-into-existence of a miserable person and the not-coming-into-existence of a happy person are both undesirable outcomes. Bringing a miserable person into existence is a case of *doing*; not bringing a happy person into existence is a case of *not doing*. Thus, assuming that there is a moral asymmetry between doing and not doing, it follows that to bring a miserable person into existence is worse than not to bring a happy person into existence.

Again, however, this view will have to take an unacceptably strong form in order to imply that it is not wrong, other things being equal, to fail to bring into existence a person whose life would be worth living. It would have to assert that, except in the case of special obligations, such as those derived through promising, we cannot be held responsible for what happens or fails to happen as a result of our not doing something. Since this view implies that it cannot be wrong, other things being equal, to fail to prevent a person from being harmed, I shall assume that it is unacceptable.

A third and seemingly more plausible way of denying that there is a moral reason to benefit people by bringing them into existence is to appeal to the principle that an act cannot be wrong unless there is or would be someone for whom that act would be bad, or worse. Call this principle the "Complainant Requirement." Since we are assuming that to bring a miserable person into existence would be bad for that person, the Complainant Requirement is compatible with the claim that it would be wrong to bring a miserable person into existence. Moreover, if we assume that it is worse for a person to fail to receive a benefit, then the Complainant Requirement is also compatible with the view that there is a moral reason to benefit people. But it also implies that it cannot be wrong not to bring into existence a person whose life would be worth living, even though to bring him into existence would be to benefit him. For, if we do not bring him into existence, there will be no one for whom that will be worse (unless, of course, it would be worse for some existing person). In short, the Complainant Requirement takes account of the important fact that, if we fail to bring a potentially happy person into existence, that cannot be bad for someone who never in fact exists.

In spite of its apparent plausibility, the Complainant Requirement is unacceptable. To see why, consider the following case.²³ Suppose that we are confronted with a choice between two social policies. One policy (the "short-term policy")

would provide certain marginal benefits for existing people, but would also cause a decline in the quality of life in the further future. This decline would not be so severe that people's lives would not be worth living, but the quality of life would be significantly lower than it would have been had we adopted the other policy (the "long-term policy") instead. The long-term policy would provide no benefits for existing people, but it would allow us to sustain a high quality of life indefinitely. It seems clear that it would be wrong to adopt the short-term policy.

The Complainant Requirement, however, implies that it would not be wrong to adopt the short-term policy. Recall the earlier claim that the identity of a person depends on the timing of his conception. Since the implementation of the short-term policy would have widespread effects on the details of people's lives, it would affect who would subsequently exist. After a certain time, there would be no one in existence who would have existed had the short-term policy not been adopted. Thus the people who would exist in the further future when the quality of life has declined would not have existed had the short-term policy not been adopted. Since their lives would be worth living, the adoption of the short-term policy cannot be worse *for them*. Since it would be better for existing people, the short-term policy is not worse for anyone who ever lives. So, according to the Complainant Requirement, it cannot be wrong to adopt the short-term policy. If we believe it *would* be wrong to adopt the short-term policy, then we must reject the Complainant Requirement.

A final possibility might be to appeal to some theory of rights. While a case can be made for thinking that one can have violated a person's rights by bringing him into a predictably miserable existence, it cannot be claimed that one would violate a person's rights by *failing* to bring him into existence. But, while the appeal to rights excludes the possibility that one could have a duty, based on a respect for people's rights, to bring into existence a person whose life would be worth living, it does not exclude the possibility that there might be some *other* moral reason for bringing him into existence. To support the conclusion that it cannot be wrong, other things being equal, not to bring such a person into existence, one would need the further claim that an act cannot be wrong unless it violates a right, which is absurd.

I know of no other way in which we can accept that to be brought into existence can be a benefit and at the same time deny that there is a moral reason to bring a person into existence if his life would be worth living. By default, therefore, I think we must accept that there is a moral reason to bring a person into existence if his life could be expected to be worth living. There is a principle of beneficence that requires us, if other things are equal, to benefit people by bringing them into existence. While this conclusion may initially seem counterintuitive, it draws further support from the fact that it helps to explain the widespread conviction that it is morally imperative to ensure the existence of future generations. Our moral reason to ensure the existence of future generations is at least in part a moral reason to provide, or not to prevent, the enormous benefits of life for the enormous number of people who might exist in the indefinite future.

This conclusion does not, however, fully account for the common belief that it is of the utmost moral importance to ensure the existence of future generations. To see this, let us compare two choices. The first is the choice between human extinction and the perpetuation of the human race. The second is a hypothetical choice between perpetuating the human race only on earth and perpetuating it both on earth and on some other planet. In this second choice, the alternative that involves populating another planet would, we may suppose, roughly double the number of people

who would exist in the future. And we may suppose that life on the other planet would be of the same quality as life here. Thus the first alternative in each of these two choices would involve denying life to roughly the same number of people: extinction would deny life to a large number of people who would have benefited from being alive, but so would the failure to populate another planet. Our intuitive conviction is that the failure to ensure the survival of the species would be far worse than the failure to populate two planets rather than one, but we cannot account for this conviction by appealing to our moral reason to benefit people by bringing them into existence. If other things were equal, our principle of beneficence would imply that the failure to populate two planets rather than one would be just as bad as the failure to ensure the survival of the species.

This problem can be solved by appealing to our first two arguments for the conclusion that it is important to ensure the existence of future generations to explain why it would be worse to fail to prevent the extinction of the species than to fail to populate two planets rather than one. The failure to prevent the extinction of the species would deprive the lives of existing people of much of their meaning and would also bring the evolution of our culture to an end, while the failure to populate two planets would not. For these reasons the failure to prevent extinction would be worse than the failure to populate two planets.²⁴

There is, however, a further objection to my argument. It is an objection to the claim that a partial explanation of why it is important to ensure the existence of future generations lies in the fact that there is a moral reason to benefit people by bringing them into existence. I shall conclude by briefly stating this objection and sketching a possible solution to it.

It is natural to assume that, when we benefit a person and other things are equal, this leads to a better state of affairs. Thus we could in principle continue to improve a state of affairs, *ceteris paribus*, just by increasing the number of lives worth living. Consider a world in which everyone has a life worth living, but of a relatively low quality. This world would become better and better the more people it contained, other things being equal. Thus, if it were sufficiently populous, it could in principle be better than *any* world with a finite number of lives, all of which would be well worth living. This has, for obvious reasons, been called "the Repugnant Conclusion."²⁵

As long as we assume that there is value in increasing the number of lives worth living, we will be threatened with the Repugnant Conclusion. But my claim that the expectation that a person would have a life worth living provides a moral reason for bringing him into existence does not *necessarily* imply this conclusion. My earlier argument does not imply that our moral reason for benefiting people by bringing them into existence must be of a certain strength, or must override other conflicting reasons for action. It is therefore compatible with the following view.

Quality of life and quantity of life are separate values. There is no compelling reason for thinking that trade-offs between them must always be arranged so as to maximize total utility. There is, in fact, *no* objectively correct set of trade-offs between the two values. But the following principles for the determination of trade-offs seem plausible, and they allow us to avoid the Repugnant Conclusion. First, increasing or preserving the quality of life in general matters more than increasing the number of lives. The force of this claim is that it always takes an *increase* in the total utility derived through increasing the number of lives to make up for a decline in the quality of life. Second, as the quality of life gets lower, it takes an increas-

ingly larger gain in the total utility derived through increasing the number of lives to make up for a fixed decline in the quality of life. Finally, once the quality of life drops to a certain point, there is *no* gain in the total utility derived through increasing the number of lives which could make up for a further decline in the quality of life, even though there is still value in increasing the number of lives worth living. In short, at this point the value of preserving the quality of life and the value of increasing the number of lives worth living become incommensurable. It follows that a world in which the quality of life is below the threshold at which incommensurability begins cannot be better than a world in which the quality of life is high. This will be true even if the world with the lower quality of life has a vastly greater total utility.

This is only the crudest sketch of how I think the Repugnant Conclusion might best be avoided. The further elaboration of this view will have to be the work of another essay.

* * *

As I mentioned earlier, the argument of this essay leaves open the possibility that some policy other than the present policy of nuclear deterrence or a policy of nonnuclear defense may be superior to both. If, for example, the strategic arguments of this paper are correct, then a policy of minimal deterrence might be justified purely on Pareto grounds, in the sense that it would be better than the present policy in some respects, and worse in none. My arguments in support of the claim that a policy of nonnuclear defense would reduce the probability of nuclear war could also be cited, with relevant changes, to show that the shift to a policy of minimal deterrence would also reduce the probability of nuclear war. (It would certainly reduce the probability that nuclear war would lead to extinction.) And, provided that it would allow for a series of counterstrikes, and provided that the weapons could remain both effective and largely invulnerable to preemption, it is arguable that the shift to a policy of minimal deterrence would not weaken deterrence, and so would not increase the probability of nuclear blackmail or Soviet domination. At least as an interim policy, minimal deterrence has much to recommend it. Indeed, the shift to a policy of minimal deterrence would be an essential preliminary to the adoption of a policy of unilateral nuclear disarmament, for it would provide an important test of the desirability of unilateral nuclear disarmament. This is because the Soviets' reaction to a unilateral shift by the US to a policy of minimal deterrence would provide important evidence of what their likely response would be to the adoption by the US of a policy of unilateral nuclear disarmament.

Alternatives such as minimal deterrence need to be more carefully explored. Because it fails to consider these other alternatives, the argument of this essay is of limited significance. But it will have served an important purpose if it has at least shown that the case for nuclear deterrence cannot be regarded as unassailable. Consideration of unilateral nuclear disarmament as a viable option can no longer be excluded from "responsible" discussion in the US.²⁶

NOTES

1. Desmond Ball, "Can Nuclear War Be Controlled?" *Adelphi Papers* no. 169 (London: International Institution of Strategic Studies, 1981); and Ian Clark, *Limited Nuclear War* (Princeton: Princeton University Press, 1982).

2. Compare Bertrand Russell, "A Counter-Reply," in *Ethics and Metaethics: Readings in Ethical Philosophy*, ed. Raziel Abelson (New York: St. Martin's Press, 1963), pp. 17 - 72.
3. Michael Howard, "On Fighting a Nuclear War," *International Security* 5, no. 4 (1981): 14.
4. Jonathan Schell, *The Fate of the Earth* (New York: Alfred A. Knopf; London: Jonathan Cape, 1982).
5. Noam Chomsky and Edward S. Herman, *The Washington Connection and Third World Fascism* (Boston: South End Press, 1979); Edward S. Herman, *The Real Terror Network* (Boston: South End Press, 1982); and Jenny Pearce, *Under the Eagle: U.S. Intervention in Central America and the Caribbean* (London: Latin American Bureau, 1982).
6. Jefferson McMahan, *British Nuclear Weapons: For and Against* (London: Junction Books, 1981).
7. *New York Times*, 9 March 1983.
8. E.P. Thompson, *Zero Option* (London: Merlin Press, 1982), p. 2. Published in the U.S. as *Beyond the Cold War*.
9. These two objections were suggested to me by Russell Hardin.
10. For further discussion of the problem of nuclear blackmail, see Jefferson McMahan, "Nuclear Blackmail," in *Dangers of Deterrence: Philosophers on Nuclear Strategy*, ed. Nigel Blake and Kay Pole (London: Routledge and Kegan Paul, 1983).
11. McMahan, *British Nuclear Weapons*.
12. Gregory S. Kavka, "Deterrence, Utility, and Rational Choice," *Theory and Decision* 12 (1980): 46 and 50.
13. See, for example, Paul Nitze, "Deterring Our Deterrent," *Foreign Policy* 25 (1976-77).
14. There are other reasons for preferring premise 6. For example, Kavka concedes that, if the probability of nuclear war under nuclear deterrence is high, then "Disaster avoidance is a rather forlorn hope and its importance pales in comparison with the goal of disaster minimization", ("Deterrence, Utility, Rational Choice," p. 51). While Kavka thinks that this probability is low, I have argued that it is significant. If I am right, then this constitutes another reason for appealing to premise 6 rather than to the Disaster Avoidance Principle. Of course, there are conditions under which it might seem reasonable to be guided by the Disaster Avoidance Principle rather than by premise 6. For example, if, as many people believe, the probability of Soviet domination under unilateral nuclear disarmament would be close to 1, then, even if unilateral nuclear disarmament would not increase the probability of Soviet domination by significantly more than it would decrease the probability of nuclear war, it might be rational to prefer nuclear deterrence. (This case was suggested to me by Robert McKim.)
15. Here and elsewhere in this essay, "he" and "his" should be understood to mean "he or she" and "his or her."
16. See, for example, Jonathan Glover, *Causing Death and Saving Lives* (Harmondsworth: Penguin Books, 1977), pp. 69-70. Like Glover, Schell (in *The Fate of the Earth*) tries to explain why it is morally important to ensure the existence of future generations; but since the crucial assumptions in his explana-

tion are both controversial and undefended, his explanation cannot be considered a justification.

17. Jonathan Bennett, "On Maximizing Happiness," in *Obligations to Future Generations*, ed. R.I. Sikora and Brian Barry (Philadelphia: Temple University Press, 1978); Joel Feinberg, "The Rights of Animals and Unborn Generations," in *Philosophy and Environmental Crisis*, ed. William Blackstone (Athens: University of Georgia Press, 1974); Jan Narveson, "Future People and Us," in Sikora and Barry, eds., *Obligations to Future Generations*; and Peter Singer, "A Utilitarian Population Principle," in *Ethics and Population*, ed. Michael Bayles (Cambridge, MA: Schenkman, 1976).
18. David Watt, "Trying to Balance on the World's Nuclear Tightrope," *The London Times*, 7 January 1983.
19. Derek Parfit, "Future Generations: Further Problems," *Philosophy and Public Affairs* 11, no 2 (1982), sect. 1.
20. This point is eloquently made by Schell, *The Fate of the Earth*, in chapter 2. I leave aside the question whether it would be worse for existing people to be the final generation even if they were not aware that they were.
21. For a rebuttal of certain objections to the claim that to be brought into existence can be a benefit, see Jefferson McMahan, "Problems of Population Theory," *Ethics* 92, no. 1 (1981): 104-9. In that paper I also argue that people benefit from being brought into existence to the full extent of their lifelong balance of utility.
22. See James Griffin, "Is Unhappiness Morally More Important than Happiness?" *Philosophical Quarterly* 29 (1979).
23. I owe this example to Derek Parfit.
24. In a review of Schell's book in the *London Review of Books* (1-14 July 1982), I unjustly criticized Schell for being unable to solve this problem. The ideas on which the solution I have proposed is based are implicit in Schell's narrative.
25. Parfit, "Future Generations."
26. This essay was written in the spring of 1983. At that time the theory of the "nuclear winter" had not been developed. Since then, however, the theory has gained increasing credibility, and the scientific community now believes that the probability that nuclear war would lead to human extinction is much greater than has hitherto been suspected. For obvious reasons, therefore, the nuclear winter findings have strengthened the argument of this essay.