

VII

Against the Social Discount Rate

TYLER COWEN AND
DEREK PARFIT

When deciding how to use resources, or to protect the environment, or when selecting other policies with long-term consequences, governments and their advisers often use a social discount rate. With such a rate, possible costs and benefits are assumed to be less important if they would come further in the future. Such a discount rate applies not only to the costs and benefits that will later come to existing people, but also across the lives of all future generations. In this chapter we consider the arguments for such a discount rate and conclude that they are not sound.¹

The Effect of Discount Rates

There are two primary methods of choosing a social discount rate. The first method uses the marginal real rate of return on private capital as a proxy for the opportunity cost of postponed consumption. One well-known study estimated such rates of returns at 12.41 percent.² After adjusting these returns for risk premiums, discount rates of between 5 and 10 percent are typically generated. The second method estimates the social rate of time preference by examining the real rate of return on the almost riskless obligations of the U.S. Treasury. This procedure usually generates discount rates between 1 and 2 percent.

If we are considering the further future, the choice of a

We wish to thank Anihai Glazer, Gregory Kavka, Daniel Klein, Randall Kroszner, Alan Nelson, and Kenneth Small for useful comments and discussions.

Table 7.1
Estimated Number of Future Benefits Equal to One
Present Benefit Based on Different Discount Rates

Years in the Future	1%	3%	5%	10%
30	1.3	2.4	4.3	17.4
50	1.6	4.3	11.4	117.3
100	2.7	19.2	131.5	13,780.6
500	144.7	2,621,877.2	39,323,261,827	4.96×10^{20}

discount rate has a significant effect on our evaluation of costs and benefits. Table 7.1 shows how many future benefits are, at various discount rates, worth as much as one present benefit.

Suppose the benefits in question are lives saved. According to a social discount rate, a single present life may be worth more than one million lives in the future. With a rate of 1 percent, these million lives must be far in the future: nearly 1,400 years. With a rate of 10 percent, the distance need only be 145 years.

Why should costs and benefits receive less weight, simply because they are further in the future? When the future comes, these benefits and costs will be no less real. Imagine finding out that you, having just reached your twenty-first birthday, must soon die of cancer because one evening Cleopatra wanted an extra helping of dessert. How could this be justified? Many arguments have been proposed.

Non-economic Arguments *The Argument from Democracy*

Many people care less about the further future. Some writers claim that, if this is true of most of the adult citizens of some democratic country, this country's government ought to employ a social discount rate. If its electorate does care less about the further future, a democratic government ought to do so as well. Failure to do so would be paternalistic, or authoritarian.

To assess this argument, we must distinguish two questions: (1) As a community, may we use a social discount rate? Are we morally justified in being less concerned about the more remote effects of our social policies, at some rate of n percent per year? (2) If most of our community answer yes to question (1), ought our government to override this ma-

majority view? The Argument from Democracy applies only to question (2).³ To question (1), which is our concern, the argument is irrelevant.

The point might be put like this. A democrat believes in certain constitutional arrangements. These provide his or her answer to question (2). How could a commitment to democracy give the democrat an answer to question (1)? Only if he or she assumes that what the majority want, or believe to be right, must be right. But no sensible democrat assumes this. Suppose that some majority want to wage an aggressive war, or care nothing about the slaughter of innocent aliens. This would not show that they are right not to care. In the same way, even if most of us do care less about the more remote effects of our social policies, and believe such lesser concern to be morally justified, this cannot show that it is justified. Whatever most of us want or believe, this moral question remains open.

It may be objected: "In some cases, this is not a moral question. Suppose that, in some referendum, we vote for a social policy that will affect only ourselves. And suppose that, because we care less about what will happen to us later, we vote for a policy that will bring us benefits now at the cost of greater burdens later. This policy is against our interests. But since this policy will affect only ourselves, we cannot be acting wrongly in voting for it. We can at most be acting irrationally."

On the assumptions that most of us accept, such claims would provide some defense of the social discount rate. But the defense seldom applies. Most social policies will affect our children, as well as ourselves. If some policy would be against the interests of our children, this could be enough to make it wrong. Similar remarks apply to the interests of those people who are not yet born. It is a moral question how much weight we ought to give to the interests of these people. When those affected have no vote, the appeal to democracy provides no answer.

The Argument from Probability

It is often claimed that we should discount more remote effects because they are less likely to occur.⁴ This argument also confuses two questions:

1. When a prediction applies to the further future, is it less likely to be correct?
2. If some prediction is correct, may we give it less weight because it applies to the further future?

The answer to (1) is often yes. But this provides no argument for answering yes to (2). Suppose we are deciding whether to cease or

increase our use of nuclear energy. We are considering possible accidents from the disposal of nuclear wastes, with estimates of predicted deaths from escaped radiation. In a small accident, such deaths might all remain statistical, in the sense that we would never know which particular deaths this accident had caused. When considering possible accidents, we must think far into the future, since some nuclear wastes remain radioactive for thousands of years. According to a social discount rate of 5 percent, one statistical death next year counts for more than a billion deaths in four hundred years. Compared with causing the single death, it is morally less important if our chosen policy causes the billion deaths. This conclusion is outrageous. The billion people would be killed further in the future. But this cannot justify the claim that, compared with killing the single person, we would be acting less badly if instead we killed a billion people.

The Argument from Probability does not lead to this conclusion. It could at most lead to a different conclusion. We know that if radiation escapes next year, we will have no adequate defense. We may believe that, over the next four centuries, some kind of countermeasure will be invented, or some cure. We may thus believe that if radiation escapes in four hundred years, it will then be much less likely to cause deaths. If we are *very* optimistic, we may think this a billion times less likely. This would be a different reason for discounting, by a factor of a billion, deaths in four hundred years. We would not be making the outrageous claim that if we do cause such deaths, each of these deaths matters a billion times less than a death next year. We would instead be claiming that these more remote deaths are a billion times less likely to occur. This would be why in our view we need hardly be concerned about the escape of radiation in four hundred years. If we were right to claim that such deaths are a billion times less likely, that conclusion would be justified. Deaths that do not occur, whether now or in four hundred years, do not matter.

This example illustrates a general point. We ought to discount those predictions that are more likely to be false. Call this a Probabilistic Discount Rate. Predictions about the further future are more likely to be false. So the two kinds of discount rate, temporal and probabilistic, roughly correlate. But they are quite different. It is therefore a mistake to discount for time *rather than* for probability. One objection is that this misstates our moral view. It makes us claim not that more remote bad consequences are less likely, but that they are less important. This is not our real view.

Another objection is that the two discount rates do not always coincide. Predictions about the further future do not decrease in certainty at some constant rate of n percent per year. Indeed, when applied to the further future, many predictions are *more* likely to be true. (Consider the predictions that some policy will have changed, or that certain resources will have been exhausted.) If we discount for time rather than probability we may thus be led to what, even on our own assumptions, are the wrong conclusions.

The Argument That Our Successors Will Be Better Off

It is sometimes claimed that we should discount effects on future generations because they will be better off than we are.⁵ When benefits or costs come to people who are better off, there are good reasons for giving them less weight. If we measure these benefits and costs in terms of resources, we can appeal to a nonutilitarian distributive principle. Benefits that are *equally* great, when received by people who are better off, may be plausibly claimed to have less moral importance.

These two arguments, though good, do not justify a social discount rate. The ground for discounting these future benefits is not that they come further in the future, but that they will come to people who will be better off. Once again, we should say what we mean. And the correlation is again imperfect. Some of our successors will not be better off than we are now. When applied to these people, the arguments just given fail to apply.

The Argument from Excessive Sacrifice and from Equality

By bearing costs now, we can give our successors greater benefits. For example, if we reinvest, rather than consume, the resulting benefits may in the end be greater. If we believe that our aim should be to maximize the total sum of benefits, and we give equal weight to benefits in the further future, the optimal rate of saving could be very high. We may seem morally required to choose policies that would impose great sacrifice on the present generation.

If these requirements seem to us excessive, we may again be led, in an attempt to avoid them, to discount effects in the further future. A typical statement runs: We clearly need a discount rate for theoretical reasons. Otherwise any small increase in benefits that extends far into the future might demand any amount of sacrifice in the present, because in time the benefits would outweigh the cost.

The same objections apply. If this is why we adopt a social discount

rate, we shall be misstating what we believe. Our belief is not that the importance of future benefits steadily declines. It is rather that no generation can be morally required to make more than certain kinds of sacrifice for the sake of future generations. And this is part of a more general view, which has nothing to do with time. On this view, no one is required to make great sacrifices merely to benefit others. If this is what we believe, this is what should influence our decisions.

If instead we express our view by adopting a social discount rate, we can be led needlessly to implausible conclusions. Suppose that, at the *same* present cost, we could prevent either a minor catastrophe in the nearer future, or a major catastrophe in the further future. Since preventing the major catastrophe would involve no extra cost, the Argument from Excessive Sacrifice fails to apply. But if we take that argument to justify a discount rate, we shall be led to conclude that the greater catastrophe is less worth preventing.

A closely related argument appeals to the claims of equality. If we aim for the greatest net sum of benefits over time, this may require a very unequal distribution between different generations. We may wish to deny that there ought to be such inequality. And we can avoid this conclusion, in some cases, if we discount later benefits. But, as Rawls points out, this is the wrong way to avoid this conclusion.⁶ If we believe that such inequality would be unjust, we should not simply aim for the greatest net sum of benefits. We should have a second moral aim: that these benefits be fairly shared between different generations. To our principle of utility we should add a principle about fair distribution. This more accurately states our real view. And it removes our reason for discounting later benefits.

The Argument from Special Relations

Some utilitarians claim that each person should give equal weight to the interests of everyone. This is not what most people believe. According to commonsense morality, we ought to give some weight to the interests of strangers. But there are certain people to whom we either may or should give some kinds of priority. Thus we are morally permitted to give some kinds of priority to our own interests. And there are certain people to whose interests we ought to give some kinds of priority. These are the people to whom we stand in certain special relations. Thus each person ought to give some kinds of priority to the interests of his or her children, parents, pupils, patients, constituents, and fellow citizens.

Such a view naturally applies to the effects of our acts on future gen-

erations. Our immediate successors will be our own children. According to common sense, we ought to give to their welfare special weight. And we have similar, if weaker, obligations to our children's children. Similar claims seem plausible at the community level. Most people believe that their government ought to be especially concerned about the interests of its own citizens. It would be natural to claim that it ought to be especially concerned about the future children of its citizens, and, to a lesser degree, about their grandchildren.

Such claims might support a new kind of discount rate. We would be discounting here not for time itself, but for degrees of kinship. But at least these two relations cannot radically diverge. Our grandchildren cannot all be born before all our children. Since the correlation is, here, more secure, we might be tempted to employ a standard discount rate. As before, this would not be justified. For one thing, on any discount rate, more remote effects always count for less. But a discount rate with respect to kinship should at some point cease to apply.⁷ We ought to give some weight to the effects of our acts on mere strangers. We ought not to give less weight to effects on our own descendants.

Nor should such a discount rate apply to all kinds of effects. Consider this comparison. Perhaps the United States government ought in general to give priority to the welfare of its own citizens. But this does not apply to the infliction of grave harms. Suppose this government decided to resume atmospheric nuclear tests. If it predicts that the resulting fallout would cause several deaths, should it discount the deaths of aliens? Should it therefore move these tests from Nevada to the South Pacific, so that those killed would not be Americans? It seems clear that, in such a case, the special relations make no moral difference. We should take the same view about the harms that we may impose on our remote successors.

Economic Arguments

We now examine the arguments for discounting that are most popular among economists: the appeal to opportunity costs, and to time preference. Unlike the arguments discussed so far, these arguments do appeal to considerations that are essentially about time. If they were sound, these arguments would provide direct support for the social discount rate. And, unlike the previous arguments, they would provide some guidance on how steep the discount rate should be.⁸

The Argument from Opportunity Costs

It is sometimes better to receive a benefit earlier, since this benefit can then be used to produce further benefits. An investment that yields a return next year will be worth more than the same return arriving in ten years if the earlier return can be profitably reinvested over these ten years. When we add in the extra benefits from this reinvestment, the total sum of benefits will be greater. A similar argument covers certain kinds of cost. The delaying of some benefits involves opportunity costs.

In the language of economics, this argument asserts that the rate of discount should be determined by the marginal rate of transformation between goods today and goods tomorrow in the productive sector of the economy—that is, the marginal product of capital. Positive discounting is justified by an appeal to the positive marginal productivity of capital.

The Argument from Opportunity Costs errs in taking the marginal productivity of capital as exogenous to other social decisions. In fact, the marginal productivity of capital depends on other social decisions, most notably the community's rate of savings. If the rate of savings is determined by how we discount the future, the productivity of capital cannot be invoked as an independent determinant of the discount rate. Instead, the choice of discount rate determines the marginal productivity of capital.

Under the assumptions usually employed by economists, additional increments of capital (additional savings) lower capital's marginal product. Each successive unit of capital invested creates less additional future output than the preceding unit. Consider investments that do not yield returns until the next generation. As long as the marginal rate of return on capital exceeds the intergenerational rate of discount, additional capital should be saved and invested. For instance, if capital yields 5 percent and we are willing to trade off present for future consumption at 4 percent, additional capital should be invested because future returns exceed our rate of discount. Optimality is achieved only when the rate of return on capital equals the discount rate applied to future consumption.

A zero rate of intergenerational discount thus implies the accumulation of capital for intergenerational investments until the marginal product of such capital equals zero. In this case, the marginal rate of return on capital can no longer be used to justify a positive rate of intergenerational discount. At the optimum point suggested by the moral principle of zero discounting of consumption streams, the marginal product of capital is also zero.

It may not be possible, however, for the marginal rate of return on capital to reach zero. The economists' assumption of diminishing marginal productivity could be wrong. Or some other constraint may prevent attainment of a zero marginal rate of return on capital.⁹ In this case capital will yield a positive marginal rate of return, whether we like it or not. The Argument from Opportunity Costs may then apply.

Even here, however, the argument fails. Although certain opportunity costs do increase over time, it misrepresents our moral reasoning to treat these opportunity costs in terms of a social discount rate. These costs should be considered directly. If instead we express these costs in terms of a discount rate, we can be led astray.

We may be led, for example, to confuse benefits that will be reinvested with benefits that are merely consumed. When benefits to be consumed are received later, this may involve no opportunity costs. Suppose we are deciding whether to build some airport. Since this airport would destroy a fine stretch of countryside, we would lose the benefit of enjoying this natural beauty. If we do not build the proposed airport, we and our successors would enjoy this benefit in every future year. According to a social discount rate, the benefits in later years count for much less than the benefit next year. How could an appeal to opportunity costs justify this? The benefit received next year—our enjoyment of this natural beauty—cannot be profitably reinvested.

Nor can such an argument apply to those costs that are merely "consumed." Suppose we know that a certain policy carries some risk of causing genetic deformities. The argument cannot show that a genetic deformity next year ought to count ten times as much as a deformity in twenty years. The most that could be claimed is this. We might decide, that for each child so affected, the large sum of k dollars would provide adequate compensation. If we were going to provide such compensation, the present cost of ensuring this would be much greater for a deformity caused next year. We would now have to set aside almost the full k dollars. A mere tenth of this sum, if set aside now and profitably invested, might yield in twenty years what would then be equivalent to k dollars. This provides one reason for being less concerned now about deformities in the further future. But the reason is not that such deformities matter less. The reason is that it would now cost us only a tenth as much to ensure that, when such deformities occur, we would be able to provide compensation. This is a crucial difference. Suppose we know that we will not in fact provide compensation. This might be true, for instance, if we would not be able to identify those particular genetic deformities

that our policy had caused. This removes our reason for being less concerned now about deformities in later years. If we will not pay compensation for such deformities, it becomes an irrelevant fact that, in the case of later deformities, it would have been cheaper to ensure now that we could have paid compensation. But if this fact has led us to adopt a social discount rate, we may fail to notice when it becomes irrelevant. We may be led to assume that, even when there is no compensation, deformities in twenty years matter only a tenth as much as deformities next year.

The Argument from Transformation

We consider next a variant of the appeal to opportunity costs; examination of this claim further illustrates the problems with that argument. Consider consumption units, which we can either eat now or plant in the ground for a positive return. The Argument from Transformation proceeds as follows:

1. 1.05 units in period two are better than 1 unit in period two;
 2. 1 unit in period one can be transformed into 1.05 units in period two.
- Therefore,
3. 1 unit in period one is better than 1 unit in period two.

This argument is invalid. The move from (1) and (2) to (3) confuses the two relations "can be transformed into" and "is as good as." It may be possible to transform a frog into a prince, but this does not imply that a frog who stays a frog is as good as a prince.¹⁰ Although 1 unit received in period one can be transformed (through investment) into 1.05 units consumed in period two, 1 unit consumed in period one and 1.05 units consumed in period two are mutually exclusive alternatives, just like the frog and the prince. Transformation may involve an increase, a decrease, or no change in value, but there is no presumption of any initial equivalence in value if the transformation does not take place.

The comparison between receiving a consumption unit in period one and receiving a consumption unit in period two can be broken down into mutually exclusive alternatives. If the first-period consumption unit is received and invested, we receive 1.05 units for period two; this is clearly better than initially receiving only 1 consumption unit for period two. But this dominance relation does not imply that first-period consumption (the mutually exclusive alternative to first-period investment) is also better than second-period consumption.

The confusion between the relations "can be transformed into" and "is

as good as" is not the only problem with this argument. The argument also ignores the possibility of investing the second-period consumption unit for period *three*. Although 1 consumption unit in period one can yield 1.05 units for the period-two generation, 1 unit in period two could yield 1.05 units in period three for the *next* generation.

For any investment opportunity for period *n* given by current receipt of a consumption unit, there exists an equivalent investment opportunity for period *n* + 1 given by receipt of a consumption unit one period later. Without some prior argument that consumption units should be discounted across generations, possessing this investment opportunity in period *n* cannot be considered superior to possessing the same opportunity in period *n* + 1. With a finite horizon, of course, resources in the current period give us one more option than resources in the next period for any decision concerning the *final* period. Nonetheless, current resources are more valuable than future resources only in the special case when current resources are invested and not consumed until the final period of time.

In summary, the opportunity costs determined by the marginal product of capital do indeed structure society's available options, but they do not support a social discount rate.¹¹ If marginal rates of substitution and marginal rates of transformation do not concur, the marginal rate of substitution should be given priority.

There are other versions of the appeal to opportunity costs, which we cannot consider here. But the central issue is, we believe, simple. When describing the effects of future policies, economists could describe the future benefits and costs in a way that used no discount rate. The arguments that appeal to opportunity costs could be fully stated in these temporally neutral terms. We believe that, on any important policy question, this would be a better, because less misleading, description of the alternatives. It would make it easier to reach the right decision.

The Argument from Positive Time Preference

Another traditional argument for discounting appeals to the fact that most people prefer to receive benefits sooner rather than later. In some cases this preference is clearly rational. By choosing to receive benefits earlier, we sometimes make them greater, or more certain. But many people have a *pure* time preference. They prefer benefits to come earlier even when they know that this will make them smaller; and they postpone costs or burdens even when they know that this

will make them greater. We incline to the view that this attitude is irrational.¹²

Even if this attitude is not irrational, it cannot justify an intergenerational discount rate. Perhaps individuals may rationally prefer smaller benefits, because they are in the nearer future. But this argument has no next step. Pure time preference within a single life does not imply pure time preference across different lives. Abstinence from consumption does not involve waiting when consumption is postponed across generations. Such abstention cannot be a meaningful bad for people who are not yet born; in the meantime, nobody is left waiting.

The Argument from Transitivity

It might be suggested that by appealing to the transitivity of the relations "better than" and "as good as" we can bridge the gap between time preference within lives and across lives. The argument might be presented as follows:

1. U_{Packet40}(1997, John) = U_{Packet40}(1997, Jim)
2. U_{Packet40}(1995, John) > U_{Packet40}(1997, John)
3. U_{Packet40}(1995, John) > U_{Packet40}(1997, Jim)

The number following U_{Packet} refers to a quantity of utility, and the information within the parentheses refers to the year the utility is enjoyed and the person who receives the utility. The symbol > means "preferred to," and the symbol = means "as good as."

This argument attempts to derive (1) from utilitarianism, (2) from positive time preference within a single life, and (3) from transitivity. But the conclusion does not follow. The argument is flawed because (1)–(3) contain two different and conflicting notions of utility.

There are two ways of viewing choice among alternatives. Either an individual's ordinal choices always coincide with higher numerical *cardinal* utility for this individual, or they do not.

Consider the first assumption. If the ordinal rankings determined by John's time preference coincide with the rankings offered by a comparison of numerical utilities, claim (2) requires revision. If one utility packet is preferred to another, they cannot both be represented by the same numerical magnitude.¹³ The utility packet on the left-hand side of (2) must be assigned a different cardinal number than the utility packet on the left-hand side of (1). It must be worth, not 40, but

some higher value, such as 42. (1) and (2) would thus read as follows:

- 1a. $UPacket40(1997, John) = UPacket40(1997, Jack)$
 - 2a. $UPacket42(1995, John) > UPacket40(1997, John)$
- We can then derive:
- 3a. $UPacket42(1995, John) > UPacket40(1997, Jack)$

But this conclusion does not express a positive discounting of utility. Instead, it simply states that a greater amount of utility now is preferable to a smaller amount of utility later. It does not deny the original claim (3) $UPacket40(1995, John) = UPacket40(1997, Jack)$, the claim that rejects the interpersonal discounting of utility.

Next consider the second kind of theory, where cardinal utilities may differ from a person's ordinal choices. John's receipt of the utility packet in 1995 is now valued at 40, rather than 42, as in (2a). When John prefers the utility packet in 1995 to the utility packet in 1997, our theory of utility no longer implies that the preferred packet is worth a greater sum (42) and the inferior packet is worth a lesser sum (40). Instead, we say that both utility packets are worth 40, but that John happens to prefer 40 in 1995 to 40 in 1997.

On this second view, if cardinal utilities and ordinal rankings contradict each other, cardinal utilities are given priority in the social welfare function.¹⁴ In this case, however, the preference relation expressed in (2) no longer expresses a social ranking. Although our approach to utility theory has stipulated that social rankings are derived from cardinal utilities, (2) derives a social ranking from ordinal preferences. If we consistently determine social rankings by cardinal utilities, (2) becomes an indifference relation. The correct sequence of relations would proceed thus:

- 1b. $UPacket40(1997, John) = UPacket40(1997, Jim)$
- 2b. $UPacket40(1995, John) = UPacket40(1997, John)$
- 3b. $UPacket40(1995, John) = UPacket40(1997, Jim)$

The objection to discounting is again confirmed. Transitivity does not allow us to bridge the gap between discounting within a life and discounting across different lives.

The Argument from Time Preference assumes that, because individuals sometimes act cavalierly toward their own future, we ought analogously to discount the utilities of other generations in the future.¹⁵ But if people were like Proust's characters and wished to postpone pleasures into the

future, would this imply that benefits for future generations ought to count for more than present benefits?¹⁶

Altruism and Bequests

We have examined only conflicting interests between different generations; in the above scenarios a fixed quantity of consumption units or utilities are available for distribution. In reality, parents usually possess some degree of altruism for their descendants. The explicit incorporation of altruism into a model of generations, however, does not fundamentally alter the problem or our arguments. We now examine two arguments that suggest that altruism provides grounds against a zero rate of intergenerational discount. One argument suggests that present generations have too little influence over discount rates; the other suggests that present generations acting through family relationships have so much influence over intergenerational allocations that the intergenerational rate of discount chosen in other decision contexts does not matter.

The Argument from Double Counting claims that in the presence of altruism, a zero rate of intergenerational discount places too much weight on the interests of future generations. The interests of future generations are counted as equal to the interests of the present generation because of the zero discount rate; the interests of future generations also receive additional weight through the altruism of the present generation. What this argument calls double counting, however, appears to be a proper counting of interests. If a person benefits from the rescue of his best friend, for instance, it is not double counting to consider as benefits both the value of the life saved and the value of the friendship to the other party.

The Argument from Offsetting Transfers does not challenge the normative validity of a zero rate of discount, but instead challenges its practical significance. As long as generations are linked by a series of altruistic bequest motives, it is argued, any attempt to redistribute resources across generations will be reversed by a change in voluntary intergenerational transfers.¹⁷ Although the use of a zero rate of discount may appear to increase the well-being of future generations, the present generation will offset this effect by decreasing the time, resources, and money devoted to helping descendants.

Assume, for instance, that a father plans to devote \$200,000 to his son's upbringing and education. Use of a zero intergenerational rate of discount, however, leads to a policy that takes \$10,000 from each member

of the father's generation to yield \$10,000 (plus epsilon) for each member of the son's generation. Given certain assumptions, it can be shown that the father's utility-maximizing response decreases his transfer to his son by 10,000 to \$190,000, undoing the intergenerational redistribution suggested by the zero discount rate. If this argument holds, our choice of discount rate for collective decisions may have little practical significance. Offsetting changes in voluntary intergenerational transfers would allow the present generation to apply whatever rate of discount it chooses to future generations.¹⁸

The Argument from Offsetting Transfers relies on special assumptions that do not always hold true. First, it assumes that all persons have children, as those without descendants would not be able to perform offsetting voluntary transfers. Second, *ex ante* voluntary transfers are assumed to exceed the redistributions caused by policy transfers; otherwise no sufficient offsetting adjustment would be possible. Perhaps the most vulnerable assumption, however, is that changes in voluntary intergenerational transfers are costless.

Changes in the size of transfers will not be costless if persons enjoy giving for its own sake. The model of offsetting transfers assumes that parents care only about the size of the transfers their children receive and not about the source of these transfers. According to this argument, parents do not regard a parental gift of \$200,000 as preferable to the child receiving all or part of this sum from other parties. But this surely misstates the nature of parental altruism; parents desire their child's welfare but also wish that they are the source of this welfare; that is, they enjoy giving to their children.¹⁹ The parental joy of giving implies that reductions in involuntary intergenerational transfers might have significant costs, and that such reductions might not be used to undo redistributions toward future generations.

The Argument from Offsetting Transfers corrects the idea that any desired intergenerational allocations can always be achieved simply by redistributing wealth. The argument does not, however, show that a social discount rate has no practical consequences.

We have discussed several arguments for the social discount rate. None succeeds. At most, these arguments might justify using such a rate as a crude rule of thumb. But this rule would often go astray. It may often be morally permissible to be less concerned about the more remote effects of our social policies. But this would never be because these effects are more remote. Rather it would be because they are less likely to occur,

or would be effects on people who are better off than we are, or because it would be cheaper now to ensure compensation, or it would be for one of the other reasons we have given. All these different reasons need to be stated and judged separately, on their merits. If we bundle them together in a social discount rate, we make ourselves morally blind.

Remoteness in time roughly correlates with a whole range of morally important facts. So does remoteness in space. Those to whom we have the greatest obligations, our own family, often live with us in the same building. We often live close to those to whom we have other special obligations, such as our clients, pupils, or patients. Most of our fellow citizens live closer to us than most aliens. But no one suggests that, because there are such correlations, we should adopt a spatial discount rate. No one thinks that we would be morally justified if we cared less about the long-range effects of our acts, at some rate of *n* percent per yard. The temporal discount rate is, we believe, as little justified.

When the other arguments do not apply, we ought to be equally concerned about the predictable effects of our acts whether these will occur in one, or a hundred, or a thousand years. This has great importance. Some effects are predictable even in the distant future. Nuclear wastes may be dangerous for thousands of years. And some of our acts have permanent effects. This would be so, for instance, of the destruction of a species, or of much of our environment, or of the irreplaceable parts of our cultural heritage.

Notes

1. Several economists examine the theoretical issues behind choosing a social rate of discount, including Arrow and Kurz, *Public Investment*; Bradford, "Constraints on Government Investment Opportunities"; Mendelsohn, "Choice of Discount Rate for Public Projects"; Warr and Wright, "Isolation Paradox"; Starrett, *Foundation of Public Economics*; Lind et al., eds., *Discounting for Time and Risk in Energy Policy*; and Lind, "Shadow Price of Capital." Economists have devoted little explicit attention to the issue of intergenerational discounting. Two exceptions are Dasgupta and Heal, *Economic Theory and Exhaustible Resources*, and Mueller, "Intergenerational Justice and the Social Discount Rate." Solow, "Economics of Resources or the Resources of Economics," notes in passing that the rate of intergenerational discount should be zero. Cost-benefit studies, however, commonly use positive discount rates without regard for whether benefits and costs are distributed over different generations; see Yang, Dower, and Menefee, *Use of Economic Analysis in Valuing Natural Resource Damages*. Cowen's discussions with economists confirm the nearly unanimous acceptance of a positive intergenerational discount rate. A number of sources examine the philosophical issues

behind intergenerational discounting. See Sikora and Barry, eds., *Obligations to Future Generations*; Parfit, *Reasons and Persons*; and Broome, "Economic Value of Life." In this essay we treat intergenerational discounting as distinct from the problem of intrapersonal discounting, which we do not challenge. Intrapersonal discounting may be more a problem of rationality than of morality.

2. See Holland and Myers, "Trends in Corporate Profitability and Capital Costs."
3. See Rawls, *Theory of Justice*, 296–297.
4. See, e.g., Dasgupta, *Control of Resources*.
5. See, e.g., Mueller, "Intergenerational Justice and the Social Discount Rate."
6. See Rawls, *Theory of Justice*, 297–298.
7. Or, to avoid discontinuity, it should asymptotically approach some horizontal level that is above zero.
8. The references cited at the beginning of note 1 provide various versions of these two arguments.
9. Many investments may yield benefits both within generations and across generations. If informational imperfections prevent the separation of the intragenerational returns from the intergenerational returns, the existence of intragenerational returns, combined with positive discounting within lives, may prevent the attainment of zero rates of return on capital.
10. Nozick, *Anarchy, State, and Utopia*, 64–65, makes a similar point in a different context. If we assert that the relation "can be transformed into" implies "is as good as," we cannot explain voluntary exchange by noting that each person prefers what he or she receives to what he or she gives up.
11. Austrian capital theory as found in Boehm-Bawerk, *History and Critique of Interest Theories*, and in Fetter, *Capital, Interest and Rent*, provides an excellent analysis of the difficulties in using the productivity of capital to determine discount rates.
12. We might, however, defend a person's right to make such choices on grounds of autonomy. Parfit, *Reasons and Persons*, examines the issues surrounding positive time preference.
13. This approach to choice implies that the concept of discounting a given amount of utility is meaningless: utility is what is left over after discounting.
14. If we postulate ordinal and cardinal rankings that do not always coincide, but continue to use ordinal rankings for social rankings, we would in effect revert back to the case where utility is determined by choice. Introducing noncoincident cardinal utilities would serve only a cosmetic function, as these cardinal utilities would not be used to determine social ranking.
15. By combining positive discounting within lives and zero discounting across lives, it may be possible, paradoxically, to increase the value of a stream of benefits simply by deferring these benefits into the future. This result, however, is not disturbing if we accept the premise of positive discounting within a life. Benefits that arrive, say, in twenty years will be enjoyed predominantly by persons who are alive now; these benefits are thus subject to a positive rate of discount for the entire twenty years. In contrast, benefits that arrive in fifty years will be enjoyed predominantly by unborns and will not be discounted for the entire fifty years. Assume, for instance, that the next generation is born in thirty-five years;

benefits arriving in fifty years are thus discounted for only fifteen years and these benefits may be more valuable than the benefits for the current generation in twenty years. However, this reflects the assumption of positive time preference within single lives; making one generation wait fifteen years for their benefits would be less bad than making another generation wait twenty years for their benefits.

16. See Loewenstein, "Anticipation and the Valuation of Delayed Consumption."
17. Examining the use of offsetting transfers as a response to intergenerational redistribution was pioneered by Barro, "Are Government Bonds Net Wealth?" Warr and Wright, "Isolation Paradox," applies Barro's argument to the choice of discount rate.
18. Advocates of the Argument from Offsetting Transfers do not restrict this argument to transfers across adjacent generations. If the present generation is taxed to create benefits for persons seventeen generations later, the present generation will adjust by decreasing transfers to their children, these children will later decrease transfers to their children, and so on, and the resulting chain reaction may restore the initial intergenerational distribution. Such intergenerational linkages are modeled by Barro, "Are Government Bonds Net Wealth?" Barro's reasoning on this point, however, has been challenged by Bernheim and Bagwell, "Is Everything Neutral?"
19. How many parents, for instance, would give up their children for adoption to benevolent millionaires?

VIII

Consequentialism Implies a Zero Rate of Intergenerational Discount

TYLER COWEN

In chapter 7, Derek Parfit and I examined the arguments used to justify a positive rate of intergenerational discount; we concluded that these arguments do not succeed. Arguments based on positive time preference within lives, positive marginal productivity of capital, the uncertainty of the future, and theories of moral obligation do not suffice to rationalize positive discounting. Costs and benefits borne by future generations should be weighted on an equal par with costs and benefits borne by persons in the present.¹

Here I present an axiomatic argument for a zero intergenerational rate of discount. I use the method of social choice theory by presenting a number of simple axioms and demonstrating the conclusions that necessarily follow. If we consider the axioms reasonable, we must also accept the conclusions. The four axioms that generate a zero rate of intergenerational discount are: (1) Pareto indifference, (2) transitivity of indifference, (3) person neutrality within generations, and (4) well-defined preferences across living in different eras.

The argument for a zero rate of intergenerational discount, however, holds only to the extent that we are consequentialists. I define consequentialism as the set of moral theories that attempts to evaluate and compare outcomes. We might instead evaluate policies in terms of rights and duties. But an intergenerational rate of discount would then become difficult to define, because we would no longer be trading off changes in outcomes at the margin. My framework contains no room for rights and duties, not

because I wish to argue for or against consequentialism, but because I wish to show its implications.

I now discuss each axiom in detail. First, Pareto indifference states that two situations are equally good if the same two situations are equally good for all persons involved. In other words, we should be indifferent toward policies that harm no one and benefit no one.²

Amartya Sen, among others, has challenged the principle of Pareto indifference by considering situations with equal utilities that differ in some other morally relevant fashion. Persons might be happy, for instance, because they possess base motives, such as enjoying the misfortune of others.³ Creating happiness through base motives might not be as good, all things considered, as creating an equal amount of happiness through love. We might have ethical grounds for preferring happiness created through love over happiness created through malice, even if the two emotions produce the same amount of utility; Sen stresses the importance of what he calls "non-welfarist" values.

Examples such as Sen's are well-taken exceptions to the principle of Pareto indifference. When I use the Pareto indifference axiom, however, the different situations I am comparing do not differ obviously with respect to non-welfarist values. Persons consume a single homogenous good across different generational eras. Within this context, I simply assume that if a person is subject to an inconvenience and simultaneously compensated for this inconvenience, no one is worse off. Paretian indifference appears an entirely reasonable axiom within the context I consider.

The second axiom, the transitivity of indifference, states that if situation A is equally good as situation B, and B is equally good as situation C, A must be equally good as C. As with the first axiom, the transitivity of indifference need not always hold. Consider the Sorites Paradox. A person may be indifferent between fifty and fifty-one grains of sugar in coffee, indifferent between fifty-one and fifty-two grains, indifferent between ninety-nine and one hundred grains, and so on, but not indifferent between fifty and one hundred grains. Intransitivity arises because of fuzzy preference rankings.⁴ Transitive indifference is a bad assumption when imperceptible changes or fuzzy preference rankings are involved, but neither of these factors is present in the issues examined here.

The third axiom, person neutrality within generations, states that a benefit for one person is equally good as the same benefit for another person within the same generation or time period, *ceteris paribus*. Granting the same benefit to either Jones or Smith is equally good. Similarly,

imposing the same cost on either Jones or Smith is equally bad. The *ceteris paribus* clause rules out such factors as asymmetrically distributed altruism, for example, that other citizens might be happier when Jones gets a benefit than when Smith gets the same benefit. By treating all persons similarly within a time period, the person-neutrality axiom ensures that any resulting asymmetrical treatment of persons is due to their location in time. For the purpose of the following discussion, benefits take the form of cardinal, interpersonally comparable utility.

The fourth axiom is that persons have well-defined preferences across living in different eras. These preferences are well defined if persons would be willing to engage in permanent time travel in exchange for some amount of money or resources. Although persons may prefer living in one era rather than another, a sufficiently large income or endowment differential equalizes their utility in different eras. I also define the concept of era indifference, which applies to a person who requires no net increase or decrease in endowments to switch eras. Persons who are era indifferent do not care when they live, provided they receive the same endowment.⁵

Preferences across eras can be made operational by the following thought experiment. If cryogenic treatments could be used to freeze persons and thaw them out many years in the future (painlessly and risklessly), how much compensation would persons require to undergo such treatments? Unlike traditional uses of cryogenics, these treatments would be applied at the beginning of a person's life, and not at the end in an attempt to avoid or postpone death. Alternatively, we might consider a *de novo* thought experiment, which gathers all persons at the beginning of time and asks them to define their preferences across eras. Preferences across eras, however, do not require the operationality of any thought experiment or method of time travel; persons need only be able to rank inhabitation of different eras, even if no actual choice is available.

The axioms of Pareto indifference, transitivity, intragenerational person neutrality, and well-defined preferences across eras, taken together, imply a zero rate of intergenerational discount. That is, one additional unit of utility for the current generation is equally good, all things considered, as one additional unit of utility for the future generation. Although both the axioms and the diagrammatical proof are defined across utilities, a similar proof can be constructed across commodities or income as well.

The underlying intuition behind the proof is the following. First, the presence of well-defined preferences across eras implies that we can conduct thought experiments that switch persons and their utility endow-

ments across eras. Of course, if era indifference does not hold, offsetting compensation in terms of endowments is required to preserve social indifference. Intragenerational person neutrality implies that it does not matter which person in a given time period receives a marginal change in utility. Combining these two results with the transitivity of indifference allows the following thought experiment: switch a person (the mover) and his utility forward in time, transfer one unit of utility from the mover to another person in the mover's new generation, and switch the mover back to his initial generation, without the unit he has lost (with compensation for the era switches but not for the lost unit of utility). We are back to where we started, except that the future generation has one additional unit of utility, and the present generation has one unit less.

At each stage in this process, the resulting situation is socially as good as the preceding situation, so social indifference relations are preserved. The initial and final situations differ only with respect to whether the present or future generation receives an additional increment of utility. The derivation of intergenerational neutrality from intragenerational neutrality follows from the extension of individual preference space across eras and the transitivity of indifference relations. Neutrality across generations becomes a natural extension of the traditional consequentialist belief in person neutrality within a single generation.

The above reasoning can be reproduced diagrammatically, where the letter designates a situation or world state, the row indicates the generation or era a person belongs to, and the information in the parentheses indicates a person's identity, material endowment, and utility, in that order. Utility is a function of a person's material endowment and other nonmaterial sources of utility; in the example presented, we are concerned with allocating nonmaterial sources of utility (for example, environmental improvements) across generations. For social rankings, of course, utilities and not material endowments matter.

We wish to compare the two possible outcomes A and E:

	A	
generation 2		(John: 20, 40)
generation 1	(James: 21, 41)	
	E	
generation 2		(John: 20, 41)
generation 1	(James: 21, 40)	

Indifference across these two outcomes implies a zero rate of intergenerational discount for utility, because we would not care if an additional increment of utility was given to the current generation (generation 1) or to the future generation.

Now consider the following comparison between A and B:

	<i>A</i>	
generation 2		(John: 20, 40)
generation 1	(James: 21, 41)	

	<i>B</i>	
generation 2	(James: 21 + k , 41)	(John: 20, 40)
generation 1		

Here k may be either positive or negative and is the change in material endowments required to keep James's utility at 41 after switching him to a new era. The existence of a finite k follows from well-defined preferences across eras, and we are indifferent between outcomes A and B because of Pareto indifference.⁶ Outcome A, of course, is the outcome that results if generation 1 receives an additional unit of utility.

Now consider outcome C, which takes a unit of utility from James and gives it to John.

	<i>C</i>	
generation 2	(James: 21 + k , 40)	(John: 20, 41)
generation 1		

We are indifferent between B and C because of intragenerational person neutrality.

Now consider world state D:

	<i>D</i>	
generation 2		(John: 20, 41)
generation 1	(James: 21 + k + z , 40)	

Here z is the change in material endowments required for James to maintain a level of welfare of 40 after being shifted back to generation 1. The existence of z follows from well-defined preferences across eras, and z may be either positive or negative. Of course, if a person has preferences only across the eras themselves and does not mind the process of switching, then $k + z = 0$.

We are indifferent between C and D because of Pareto indifference.

Transitivity implies that we are indifferent between A and D through the indifference relations with B and C. Now consider world state E again:

	<i>E</i>	
generation 2		(John: 20, 41)
generation 1	(James: 21, 40)	

E is equally good as D because of Pareto indifference. That is, James does not care if he receives a given amount of well-being through material endowments or directly through utility.

The difference between A and E is the world state that results if we give an additional unit of utility to generation 2 instead of generation 1. Since D has been shown to be equally good as A, and E is equally good as D, transitivity of indifference implies that E is equally good as A, which implies a zero rate of intergenerational discount. We should not care which generation receives an additional increment of utility.⁷

The above proof is not a knockdown argument for a zero rate of discount. A non-consequentialist might simply refuse to accept any of the four axioms outlined above. Some might believe that the allocation of intergenerational resources should be determined by Kantian duties and obligations, for instance, rather than by a comparison of outcomes.

Policy analysts, however, use a consequentialist framework when performing cost-benefit studies. That is, policy analysts try to find the best outcome for society, measured by some standard such as utility, wealth, or Paretian optimality. Rejecting the axioms discussed above implies a rejection of policy analysis as a method, and not merely the rejection of a zero intergenerational rate of discount. The alternative to a zero intergenerational rate of discount, then, is not a positive rate of discount, but an unwillingness to evaluate outcomes by comparing costs and benefits.

Notes

1. See Cowen and Parfit in this volume. References to the relevant literature can be found there.
2. My examples assume that the number and the identity of persons are fixed, thus avoiding the well-known paradoxical results generated by Parfit. *Reasons and Persons*.
3. Sen, *Choice, Welfare, and Measurement*, offers a number of essays critical of Paretianism.

4. Consider another example. The fuzziness of rankings might imply an indifference relation when comparing either Chopin's piano music or Schubert's piano music to Mahler's symphonies. Even given this fuzzy indifference ranking, we might still unequivocally prefer Chopin's piano music to Schubert's.
5. Under pure era indifference, living from 1962 to 1988 is equally good as living from 2062 to 2088, given equivalent endowments. Preferences characterized by pure era indifference exhibit time preference only across intervals when individuals are living or conscious. There may be many reasonable violations of pure era indifference, however. A person may wish to have met Lord Byron or may have an aesthetic preference for belonging to a certain generation. As long as these preferences are not lexicographic, however, an indifference point across eras will still exist; some amount of compensation will induce a person to switch eras.
6. The compensation for switching eras need not actually come from another person; it must only be true that the outcome in which additional income compensates a person for moving is equally as good as the preceding world state. Compensation is treated as manna from heaven, and it does not affect the proof if compensation is not actually available. The series of world states with compensation are only midpoints used for constructing an indifference relation across the feasible world states resulting from the policy choice that gives additional utility either to generation 1 or to generation 2. Whether these midpoints are feasible does not affect the resulting indifference relation between two world states that are feasible. The following analogy illustrates why compensation need not be available for the proof to follow. We might use transitivity to conclude that if preserving the Amazon forest is equally good as irrigating the Sahara and irrigating the Sahara is equally good as preserving Antarctica, then preserving the Amazon is equally good as preserving Antarctica. If it later turns out that irrigating the Sahara is very costly and therefore impossible without manna from heaven (although still valued the same, if it could be achieved without manna), this does not break the indifference relation between preserving the Amazon and preserving Antarctica.
7. Although the above argument portrays each person as inhabiting a single well-defined generation, incorporation of overlapping generations does not affect the proof. The proof does not rule out the existence of other persons whose generations overlap or coincide. The diagrammatical demonstrations assume only that the future will contain people who do not yet exist.

Intergenerational Inequality

LARRY S. TEMKIN

Views about the moral relations between generations may be as varied, and as complex, as views about morality itself. Some see conflicting claims or interests of different generations in utilitarian terms, while others may see the issue in Aristotelian or Kantian terms. Still others may view the issue in Nietzschean, Marxian, Rawlsian, or Nozickian terms. And of course there are other views, including contractualism, relativism, and nihilism.

I am a pluralist. I believe there is a kernel of truth to many moral views, including most of those noted above. Correspondingly, I think the issue of the moral relations between generations is very complex. Indeed, to even present, much less argue for, the many strengths and weaknesses of the relevant moral views lies well beyond the scope of a single chapter. Therefore, this chapter adopts a limited strategy. It focuses on *one* moral ideal, equality, and raises various questions regarding that ideal. Unfortunately, even this limited strategy requires addressing a surprisingly broad set of issues and factors. Moreover, given this limited strategy and the nature of my inquiries, many of my results will themselves be limited or conditional. Still, I believe this chapter serves several purposes. It reveals that various answers might be given to my questions, which have important implications for the conflicting claims or interests of different generations. In addition, in settling certain questions, while raising and leaving others open, this chapter provides direction and helps lay the foundation for future inquiries regarding the complex topic of intergenerational inequality.