

Can We Use Human Judgments to Determine the Discount Rate?

Jonathan Baron¹

It has been suggested that the long-term discount rate for environmental goods should decrease at longer delays. One justification for this suggestion is that human judgments support it. This article presents an experiment showing that judgments concerning discount rates are internally inconsistent. These results point to potential problems with the use of judgments referenda for determining discount rates in cost-benefit analyses.

KEY WORDS: Discounting; judgments; biases; dynamic inconsistency

INTRODUCTION

Policy decisions made today affect people in the future who do not participate in making the decisions. Most policy decisions are made at the national level, so they can also affect people in other nations who also do not participate. In this way, national decisions and decisions about the future are similar. They involve externalities—effects on nonparticipants. Examples of policies that influence future people are those that affect future biodiversity, such as forest management, trade (in wood), population control and zoning. Examples of policies that affect other countries include investments in research (e.g., on tropical diseases or agriculture), trade, debt relief, and foreign aid. The general issue addressed in this article is how decisions of this sort should be made.

Decisions that affect the future have been debated in terms of cost-benefit analysis.⁽¹⁾ When cost-benefit analysis compares present costs with later costs, the latter must be discounted. A million dollars today will be worth quite a bit more in 25 years if it is invested or it will cost more in 25 years if it is bor-

rowed and repaid with interest. It is natural to also discount benefits, otherwise under reasonable assumptions, all beneficial projects would be postponed while society waits for its stock of money to grow ever larger.⁽²⁾ In particular, this argument assumes that society can continue to trade off money and other benefits at roughly the same rate, that is, that the benefits are fungible.

How much should the future be discounted in cost-benefit analysis? The use of current interest rates seems to discount the future too much, and use of zero discounting implies that huge amounts of money ought to be spent to prevent global warming or the loss of biodiversity.

It may help to see this problem as an example of a more general sort of trade-off, one between society (the decision makers), and others. Both kinds of policies—those that affect people in other nations and those that affect future people—involve such trade-offs. If society were completely selfish, it would ignore these effects and make decisions for its own benefit only. To be perfect utilitarians, believing that everyone counts the same, society would maximize utility for everyone involved. To judge from actual behavior, society is somewhere between perfectly selfish and perfectly utilitarian. We give others less weight than ourselves, but usually not zero weight.

¹ Department of Psychology, University of Pennsylvania, 3815 Walnut Street, Philadelphia, PA 19104-6196; baron@psych.upenn.edu.

When competition is involved, however, the weight may be negative.

Even perfect utilitarians have some reason to discount effects on people living in the future. First, the future is uncertain in that the farther away in the future events are, the ability to predict them accurately decreases. A reasonable way to approximate this loss in predictability is to assume that it results from uncertain events, such as technological changes, that could occur with equal probability at any time. Under this assumption, the expected utility of a future event declines exponentially as a function of time.

Second, marginal utility of goods is declining, and this affects the optimal distribution of material goods. People in poor countries now derive greater utility from a given amount of goods than people in rich countries. People in the future are likely to be richer than people today, so they are likely to derive less benefit from a given additional amount of goods.⁽³⁾

Beyond these considerations, though, people who make policy—and that includes all of us in our role as citizens—may weigh their own utilities more heavily than the utilities of others. The discounting of others' utilities may fall off as a function of distance. Family members may be weighed more than others, as may those who belong to the same nation, race, or ethnic group than those who do not. Likewise, we may weigh our great-grandchildren's generation more heavily than the generation of their great-grandchildren.

Pure utilitarians do not find such differential weighing to be justified. But, given that it is inevitable for decision makers, it is typically better for others if weighing is consistent according to the effects on those others. That is, society should look for the greatest benefit to others in the same category, given a fixed sacrifice of its own interest, or society should look for the smallest sacrifice required for a given benefit. Imagine two programs that will save lives of people living 100 years from now, one at the cost of \$1 million per life and the other at the cost of \$2 million. Society should put all its money into the first and none into the second, up to whatever amount society is willing to spend. Yet, if people are weighed inconsistently, money will sometimes be put into less-effective programs, thus hurting present society, others, or both.

Somewhat in this spirit, some authors have suggested that the discounting of long-term future effects should be based on values elicited from current citizens, those who must make the sacrifice in question for the benefit of others.⁽⁴⁾ Surveys of citizens'

values might be important to governments that want to satisfy public opinion without insisting that citizens get involved in all the nitty-gritty details of policy. At least, if there was a consistent number to indicate how much outcomes should be weighed as a function of their temporal distance, the losses that result from inconsistent allocation to projects that have their effects at the same time could be avoided.

Dynamic Inconsistency

One problem that arises is that every attempt to measure discount rates by asking people questions, or even by observing individual decisions, has found dynamic inconsistency. Specifically, the rate of discounting declines as the effects are farther away in the future from the time the decision is made. This means that people would make different choices concerning the same outcomes, depending upon when the decision is made.

From a utilitarian perspective, this makes no sense so long as outcomes can be assigned to specific times. The goodness, or utility, of an outcome must depend on the extent to which it achieves people's goals as expressed in their fundamental objectives, the criteria by which they evaluate a state of affairs. Utility therefore cannot simply depend on when a decision is made. Note that some utilities cannot be assigned to specific times in the way required; for example, the utility of hearing a beautiful piece of music cannot be assigned to a point in time. But this sort of example is largely irrelevant to policy decisions with effects over decades and centuries.

Consider a simple case. It is 2:59 P.M. A child has a choice of one piece of candy at 3:00 or two pieces at 4:00. Many children will take the first option. Suppose that such a child is given the same choice at 2:00 instead of 2:59. That is, the choice is still between one piece at 3:00 or two pieces at 4:00. It is easy to imagine that the child will choose the second option. In this case, the child has made a different choice solely as a function of time (putting aside for the moment the child's emotions of frustration and anticipation). This pattern of choices is *dynamically inconsistent*, that is, inconsistent over time, and is nonnormative, because the time at which the decision is made does not affect the extent to which the options achieve the child's goals. If the child's main goal is to get as much candy as possible, for example, then the second option should be chosen in both cases. Dynamic inconsistency violates the principle of "delay independence."

To avoid dynamic inconsistency, discounting

must be exponential. Researchers have repeatedly found, in both animals and people,⁽⁵⁻¹⁰⁾ that discounting is more like a hyperbola than an exponential function. Discounting is inconsistent in other ways, as well.⁽¹⁰⁻¹²⁾ The discount rate is higher when rewards are smaller, and the discount rate for losses is lower than that for gains. These results are found both with hypothetical questions and with experiments in which decisions lead to real rewards (money for people, food pellets for hungry pigeons).

Hyperbolic discounting might reflect a tendency to see events closer in time as more distinct from each other than events farther in the future. This, however, is a kind of perceptual distortion. If such distortion operates when people consider policies that affect the long-range future, then their evaluation of these policies need not reflect their attitude toward the utility or well-being of those affected. It certainly need not reflect the utilities of those affected, but it has already been acknowledged that society must live with a certain amount of partiality toward self, family, and cohort, if not co-nationality and co-race.

On the other hand, hyperbolic discounting could reflect a true opinion about how society wants to treat others. The present generation could, for example, care very much about those close to itself in time, and less about others after that, but nearly the same for others after that, no matter when they live. This would lead to policies that future people, with different utilities, would want to overturn. They would make a larger distinction than the present generation between those directly after them and those in the more distant future. Of course, people in the future will have all sorts of reasons for wanting to overturn present-day decisions.

In sum, the question addressed here is whether people's judgments about how they want to weigh the utilities of others can be taken at face value. If these judgments are the results of distortions, then they could be internally inconsistent, and therefore useless. The problem is not so much that future generations will want to undo the decisions made today; that is almost inevitable. Rather, it may turn out that judgments made today depend on how society is asked about them, on what perspective society is asked to take.

AN EXPERIMENT

To demonstrate how judgments can be inconsistent, and hence invalid as measures of utility for outcomes, three experiments were conducted. Only the last is reported, since the results of the first two were

the same and the last is more complete in the variables it examines. Subjects were asked about programs that saved lives or species, now or in the future. Questions specifically asked were

- How many units (lives or species) saved 25 years in the future were equivalent to saving 100,000 now?
- How many units (lives or species) saved 50 years in the future were equivalent to saving 100,000 now?
- How many units at 50 years were equivalent to saving 100,000 at 25 years?
- How many units at 50 years were equivalent to saving 100,000 at 25 years, if the decision were made in 25 years?

The last question is comparable to the first, since the delay in both cases is 25 years from the time the decision is being made versus at that time.

The same questions were also asked multiplying the years by two, so that time frames were 50 and 100 years instead of 25 and 50. This manipulation is designated "scale."

The questions asked allow for two consistency checks. First, responses to 0 versus 50 years, for example, should be inferable from 0 versus 25 years and 25 versus 50. In particular, by examining the ratio of the response to 100,000 units saved and calling that the "compensating ratio," then (assuming that utility is a power function of the number of units) the compensating ratio for 0 versus 50 years should be the product of the other two. This is a test of delay independence. Unlike previous studies, however, the items were next to each other and subjects could try to avoid the effect.

The second consistency check involves the question about what the judgment would be if the decision were made in 25 (or 50) years. Answers to this should be the same as the compensating ratio for 25 versus 50 years. The hypothesis is that the ratio is smaller.

The manipulation of scale also tests another hypothesis. When asking people about discounting the future for social programs, two different kinds of effects have to be considered, those that occur within the lifetime of individuals and those that affect different people who exist at different times. It might be expected that it is possible to discount at different rates, for example, caring differentially about people who exist 50 and 100 years from now. Another example is that the present generation might care more about the immediate present than 25 years hence for people who are alive at both times. This would lead to a high

discount rate for the next 25 years and a zero rate once everyone alive has died. Although such an argument runs into many technical problems,⁽¹⁾ people may think this way thereby affecting their judgments in tasks that involve discounting. They will have a higher discount rate for a smaller scale.

Three other variables were also manipulated. The first was whether the effects of the program would be felt in poor countries or rich countries, the hypothesis being that people in rich countries would see the decisions as between “helping ourselves now” and “helping others.” If so, the “others” would be discounted largely because they are not “us,” and, once this discounting occurred, further discounting for delay into the future would not occur. In this case, the discount rate for helping people in poor countries would be lower than that for helping people in rich countries.

The second variable was lives versus species. Although there was no particular hypothesis about whether the discount rate would be higher for one or the other, people might think of species as more of a world resource so that *where* the species existed (poor vs. nonpoor) would not matter.

The third variable was fungibility, the possibility of saving money now and spending it later. As noted earlier, one of the arguments for discounting benefits as well as costs is that, if money is saved now, it can be invested and the proceeds used later for similar projects. If future benefits are valued as much as present benefits, and if the proceeds from investment can be usefully spent in the future, it is always better to invest the money now and spend it in the future than to spend it now (and the same argument holds for future decisions). If the money must be spent now, though, future benefits and current benefits can be considered equal. Future benefits might thus be discounted at a lower rate when expenditures are not fungible. In each case, either subjects were reminded of the possibility of delaying expenditures or the possibility was explicitly ruled out. Previous research has found that the correlation of discount rates for health and money increases when health and money are fungible in this way.⁽¹¹⁾

Method

Sixty subjects completed a questionnaire on the World Wide Web, for \$3 each. To receive the \$3, the subject had to provide a name, address, email address, and Social Security number if residing in the United States. The questionnaire was written in JavaScript, which randomized the order of 16 conditions

separately for each subject, and carried out various checks after each answer to insure that the subjects were answering with sufficient care. Four subjects were eliminated for answers that seemed to reflect misunderstanding (and which were missed by the checks), leaving 56.

The introduction reads as follows (with some minor editing).

Species and lives

This is about two problems, species and health. The species problem is that human activity has increased the rate at which species of plants, animals (including insects) become extinct. Some estimates say that as many as 500,000 species disappear each year (out of the world's 100,000,000 species) (Ehrlich and Wilson, *Science* magazine, Aug. 16, 1991).⁽¹³⁾ This results from pollution, fishing, and using land for agriculture, roads, buildings, etc.

The health problem is that about 25,000,000 people around the world die young each year from health conditions that can be prevented or cured (out of 6,000,000,000 people and 50,000,000 total deaths). These include respiratory infections, diarrhea, lung cancer, malaria, accidents, and AIDS (Murray and Lopez, *Science*, Nov. 1, 1996).⁽¹⁴⁾

We, that is, people throughout the world, might spend money to slow down the loss of species or prevent early deaths. This study is about how we evaluate programs that do these things. In all cases, imagine that the programs are paid for by contributions from many nations based on ability to pay. Nations classified as “**poor**” pay nothing. These are, according to the U.N.: Afghanistan, . . . [countries were listed], and Zambia. Here I call other nations “**non-poor**”.

The programs are efficiently run, and as effective as they are claimed to be. The life-saving programs save, on the average, 20 years of life for each person “saved,” and the people saved range from children to older people.

Some of the programs affect only the poor nations, and others affect only the non-poor nations. For example, some health problems, like cholera, occur mostly in poor nations, and others, like automobile crashes, occur mostly in non-poor nations. And each nation has different species of plants and animals. Please assume that the poor nations will continue to be poor in the future—just as poor as they are now.

For each choice, you get two options, for example:

Option A: Save 100,000 species now.

Option B: Save ,000 species 25 years from now.

You have to enter a number that makes options A and B equally attractive. If the number of species in B were higher than the number you enter, you should prefer B. If the number of species in B were lower than what you enter, you should prefer A. You should not write LESS than 100. If it is better to save species now than later, you would have to save more later to make up for it. (But 100 is OK.)

The items also differ in when the species or lives are

saved. Programs of this sort can have delayed effects. It may seem unrealistic to say that some program will do something in exactly 50 years. But it is not unrealistic to think that some part of a program will do that. And the purpose of this study is to see how you think about timing of these benefits. So try to imagine that the statements are true.

The cost of the programs is fixed at a certain amount of money per life or per species.

The items differ in whether money saved [by not spending it now] can be invested for later use in the same kind of program. When money can be invested for future use, the interest rate on the money is 5% in real interest, and the cost of the future program is the same in current dollars. This means that you can suppose there is no inflation. When money cannot be invested, it is returned to the taxpayers in the form of reduced taxes.

In some items, you put yourself in the position of someone making the decision in the future.

There are 16 screens, each with four items. The screens may look alike, but they are all different. So please pay attention to the differences. (Some differences may not matter to you, but pay attention anyway.)

Each screen took the following form (with alternatives in brackets):

Any money you save by spending less **can [cannot] be invested** at 5% interest and used later to save **lives [species]** for the same cost [is returned to the taxpayers and **cannot be used later** to save **lives [species]**.

This decision affects only **poor [non-poor]** nations.

Option A: Save 100,000 lives [species] **now**
 Option B: Save ,000 lives [species] **25 [50] years from now**
 Option A: Save 100,000 lives [species] **now**
 Option B: Save ,000 lives [species] **50 [100] years from now**

Option A: Save 100,000 lives [species] **25 [50] years from now**
 Option B: Save ,000 lives [species] **50 [100] years from now**

Now imagine making this decision in 25 [50] years:

Option A: Save 100,000 lives [species] **then**
 Option B: Save ,000 lives [species] **25 [50] years from then**

The 16 conditions, presented in a random order, were all combination of fungible vs. non-fungible (the top item on the screen), poor vs. non-poor, species vs. lives, and 25 vs. 50 year steps. For the 25 year steps, the four items on each screen were now vs. 25, now vs. 50, 25 vs. 50, and 25 vs. 50 but from the perspective of someone in 25 years. For the 50 year steps, all the delays were double.

Table I. Mean Compensating Ratio (Based on Geometric Means) as a Function of Time Period and Time of Decision

Time period (years)	Compensating ratio	Annual interest (%)
0–25	2.30	3.40
0–50	3.78	2.70
25–50	2.31	3.40
25–50 at 25	2.50	3.74
0–50	3.21	2.36
0–100	5.54	1.73
50–100	3.07	2.27
50–100 at 50	3.51	2.54

Note: “At 25” means the decision is made at 25 years. The bottom four rows are from the screens with 50-year steps. The right column shows the annual interest inferred from the ratio.

Results

All responses were converted to logarithms before analysis; hence all means are based on geometric means, both those computed within subjects and those computed across subjects. All statistical tests were based on these geometric means, and back transformation was done only to enter numbers in Table I. Species versus lives, poor versus nonpoor, and fungible versus nonfungible had no significant effects on discount rate and did not interact with each other or with scale (25- versus 50-year steps). Further analysis collapsed across these variables. Table I shows the results, collapsing over the variables that had no effect. Several results are apparent from the analysis.

Sensitivity to Scale

Subjects were somewhat sensitive to scale. The ratios in the bottom half of the table, where the step size was 50 years, are larger than those in the top half, $t(54) = 6.72, p = 0.0000$. But the adjustment is insufficient, because the implied discount rates are lower when the step size is larger, $t(54) = 7.62, p = 0.0000$. The clearest comparison is that between the 0 to 50-year period when it is the longest interval on the screen (second row of table) and when it is half of the longest (fifth row). The ratio is larger when it is the longest, $t(54) = 4.11, p = 0.0001$. In sum, subjects seem to evaluate time intervals by comparison to what is immediately available (in this case, on the same screen). This result is consistent with the psychological-distortion explanation, rather than the idea that people prefer the current and next generation and treat other generations nearly alike.

Discount Rate and Delay

As found before by other researchers, discount rates (as interest percentages) are not constant as a function of delay. The discount rates for the longer intervals are all significantly lower than those for the shorter intervals. This result, however, is not a function of a general judgment that the discount rate should decline in the future. In particular, the discount ratio for 0–25 years is the same as (not significantly different from) that for 25–50 years—based on the choice between spending money in 25 years or 50 years—and likewise for 0–50 years versus 50–100.

The rate estimated from a decision made at a future time, however, is higher than the same decision made “now”; in an analysis of variance that included both scale sizes, $F(1, 54) = 7.58, p = 0.0080$, the interaction was also significant, the effect being larger for the larger scale, $F(1, 54) = 4.75, p = 0.0337$. Subjects seemed to think that they would discount more if they were transported into the future, but they did not in fact discount more for time intervals beginning in the future than for intervals beginning now, when they made the decision “now.”

Consistency of Discounting

Discounting from 25 to 50 years should equal the ratio of discounting from 0 to 25 years, as well as from 0 to 50 years. This hypothesis was tested by looking at the difference of the predicted versus obtained log compensation ratios. The ratio for 25–50 years predicted from 0–25 years and 0–50 years was 1.64, which was significantly lower than the obtained 25–50-year ratio of 2.31, $t(55) = 5.05, p = 0.0000$. Likewise, the ratio for 50–100 years predicted from 0–50 and 0–100 years was 1.73, which was lower than the obtained 50–100-year ratio of 3.07, $t(55) = 6.42, p = 0.0000$. This effect is easily explained in terms of insufficient sensitivity to delay. The difference between a 25-year interval (0–25) and a 50-year interval (0–50) is too small.

IMPLICATIONS

These results suggest that the finding of declining discount rates in questionnaire studies may be an artifact of judgment. As has been found in many other studies,⁽¹⁵⁾ people are insensitive to quantity. Here, they are insensitive to the amount of delay. They distinguish between present and future, but they are insufficiently sensitive to the amount of future delay. They are also insensitive to scale when this is manip-

ulated across items. They seem to make judgments primarily by comparing the cases they see. These effects lead to inconsistent judgments.

Moreover, in this study, subjects were not sensitive to fungibility, even though their implicit discount rates were usually below the 5% interest rate described.

Nor were there different discount rates between “us” and “them.” In this study, at least, the preference for the immediate present applies even though the effects are entirely on other people—strangers in foreign countries. Such a judgment may be challenged. If we do not know the people our programs affect, why should we care when they live? The general conclusion is that these judgments cannot be taken to be judgments about the differential weight of consequences that occur at different times.

Fairness

Such inconsistency might also affect judgments about the allocation of resources among people living now. When allocating resources, various heuristics are applied, such as equality or deservedness. These heuristics can lead to inconsistency. For example, when given a fixed budget for cancer screening, people would rather offer a less-effective screening test to 100% of a population, thereby saving 1,000 lives, than a more effective test to 50% of the population, thereby saving 1,100. When both percentages are reduced by a constant proportion or a constant amount, however, the number of people preferring the less-effective screening test was reduced.⁽¹⁶⁾ This effect can be induced as a framing effect, simply by doubling the size of the “population” at issue, which makes 100% versus 50% into 50% versus 25%.⁽¹⁷⁾ Even though the people affected are the same, the judgment changes. When it is impossible to treat “everyone” equally by giving them all the test, people are no longer willing to sacrifice lives for the sake of equality. Yet “everyone” depends on which group is being discussed.

Of course, the analogy between this sort of result and the kind of inconsistencies shown in the present experiment is loose. When society makes judgments about the future, it tends to think in terms of distance, with no sharp boundary. When considering fairness to groups, though, there tend to be sharper boundaries, such as that between co-nationals and foreigners. What is common to the two scenarios, though, is that the judgments made from different perspectives disagree.

If the concern was only with outcomes for individuals, ignoring when or where they live (or their

race or gender, for that matter), judgments would be more consistent. This can still be done while weighing ourselves and our families more than others.

Democracy

Kopp and Portney⁽⁴⁾ recently proposed that intergenerational decision making be done through a mock referendum, in which a random sample of citizens were given all the information available about the costs and benefits of alternative policies, such as those concerning global warming. Then the respondents would vote on various proposals, such as carbon taxes of various amounts. One argument for this proposal was that the issues were so complex and uncertain as to make cost-benefit analysis, with a fixed discount rate, less useful as a tool. The second main argument was that, if the people in such a referendum rejected some proposal, under such maximum-information conditions, it would not pass the legislature anyway.

One difficulty with this proposal is that it does not really solve the problems of complexity and uncertainty. It just pushes off this problem onto the respondents, as if they were some sort of “black box” that could deal with complexity better than any policy analyst. Some respondents, however, get frustrated wading through lots of detailed information when what they really want is the results of a thorough policy analysis, even just a best guess, as inaccurate as it may be. For example, if I am asked my opinion of the Kyoto Protocol, I do not want to read it, plus all the supporting technical documentation. Even less do I want to read a watered-down version of it meant for public consumption. What I would like is the summary opinion of experts about whether the benefits exceed the costs. The referendum idea does not circumvent the need for policy analysis, for those respondents who might care about such analysis.

Another argument for a referendum is that any democratic government that ignores the final expression of public preferences, whatever their basis, will not stay in power for long. But this is surely an exaggeration. The public’s preferences may be transitory.⁽¹⁸⁾ Or citizens may come to recognize that they lack sufficient expertise to make the final decisions. The U.S. Federal Reserve, for example, does not do surveys to decide what interest and inflation rates people would like. People can, up to a point, understand that their government ought to ignore their immediate preferences. Breyer⁽¹⁹⁾ argues that this can happen even in the controversial area of risk regulation, although his proposals have not attracted much

interest. An enlightened government could try to maximize true utility *and* explain what it is doing, so that it can maintain public acceptance. This idea is fully consistent with democratic governance.

CONCLUSION

What should governments do, and what should citizens expect of them? The general type of decision at issue is one in which present citizens must reduce their current well-being for the sake of other people, either those living in the future (within the same nation or not), or those from other nations.

One thing governments can do is try to improve their capacity for understanding public preferences. One strategy for doing this is to present respondents with their own inconsistencies and then ask the respondents to resolve the inconsistencies. People seem not to object to this procedure, and, when it is done, other inconsistencies are reduced, in addition to the ones that the respondents resolved.^(20,21)

In principle, such improved methods of value measurement could determine how much altruism people have, that is, how much they are willing to weigh the interests of nonnationals and future people, relative to their own. Extensive application of such methods might involve a kind of education in which people are challenged not only to be consistent over time but also in their treatment of future nationals and nonnationals. (No such inconsistency was observed here, but it might be observed elsewhere.) Indeed, people might be challenged on their preference for current co-nationals—other than themselves and their families—over others, a preference that surely exists yet is arguably as arbitrary as preference for one’s own race or sex.^(22,23)

In the meantime, however, while society waits for value measurement and people’s values to become more consistent and less arbitrary, governments might at least try—with the consent and knowledge of their citizens—to bring some internal consistency into their policies that affect nonnationals and future people. It isn’t necessary to assign a monetary value on lives in order to do this. Instead, simple cost-effectiveness analysis can be applied. Society can look for cases in which a lot of money is being spent for little benefit to others and other cases where much less money could be spent to do more good, and money could be transferred from the former to the latter. Nobody can make these judgments except experts who can figure the costs and benefits. If governments would adopt the policy of relying on such

experts, it might be able to earn the trust of citizens, just as it has done in the matter of setting interest rates.⁽¹⁹⁾

ACKNOWLEDGMENT

Research for this article was supported by National Science Foundation Grant No. SES98-76469.

REFERENCES

1. Portney, P. R., & Weyant, J. P. (Eds.). (1999). *Discounting and intergenerational equity*. Washington, DC: Resources for the Future.
2. Keeler, E., & Cretin, S. (1983). Discounting of life-saving and other non-monetary effects. *Management Science*, 29, 300–306.
3. Schelling, T. C. (1996). Global decisions for the very long term: Intergenerational and international discounting. In R. J. Zeckhauser, R. L. Keeney, & J. K. Sebenius (Eds.), *Wise choices: Decisions, games, and negotiations* (pp. 152–166). Boston: Harvard Business School Press.
4. Kopp, R. J., & Portney, P. R. (1999). Mock referenda for intergenerational decisionmaking. In P. R. Portney & J. P. Weyant (Eds.), *Discounting and intergenerational equity* (pp. 87–98). Washington, DC: Resources for the Future.
5. Ainslie, G., & Haendel, V. (1983). The motives of the will. In E. Gottheil, K. A. Druley, T. E. Skoloda, & H. M. Waxman (Eds.), *Etiologic aspects of alcohol and drug abuse* (pp. 119–140). Springfield, IL: Thomas.
6. Chung, S. H., & Herrnstein, R. J. (1967). Choice and delay of reinforcement. *Journal of the Experimental Analysis of Behavior*, 10, 67–74.
7. Solnick, J. V., Kanneberg, C. H., Eckerman, D. A., & Waller, M. B. (1980). An experimental analysis of impulsivity and impulse control in humans. *Learning and Motivation*, 11, 61–77.
8. Cropper, M. L., Aydede, S. K., & Portney, P. R. (1992). Rates of time preference for saving lives. *American Economic Review*, 82, 469–472.
9. Atherton, E., & French, S. (1999). Valuing the future: An application of hyperbolic discounting with time weights. *Risk Decision and Policy*, 4, 17–29.
10. Kirby, K. N. (1997). Bidding on the future: Evidence against normative discounting of delayed rewards. *Journal of Experimental Psychology: General*, 126, 54–70.
11. Chapman, G. C. (1998). Sooner or later: The psychology of intertemporal choice. In D. L. Medin (Ed.), *The psychology of learning and motivation* (Vol. 38, pp. 83–113). New York: Academic Press.
12. Loewenstein, G., & Prelec, D. (1993). Preferences for sequences of outcomes. *Psychological Review*, 100, 91–108.
13. Ehrlich, P., & Wilson E. O. (1991). Biodiversity studies: Science and policy. *Science*, 253, 758–762.
14. Murray, C. J., & Lopez, A. D. (1996). Evidence-based health policy—lessons from the Global Burden of Disease Study. *Science*, 274, 740–743.
15. Baron, J. (1997). Biases in the quantitative measurement of values for public decisions. *Psychological Bulletin*, 122, 72–88.
16. Ubel, P. A., Baron, J., Nash, B., & Asch, D. A. (2000). Are preferences for equity over efficiency in health care allocation “all or nothing”? *Medical Care*, 38, 366–373.
17. Ubel, P. A., Baron, J., & Asch, D. A. (in press). Preference for equity as a framing effect. *Medical Decision Making*.
18. Loewenstein, G., & Mather, J. (1990). Dynamic processes in risk perception. *Journal of Risk and Uncertainty*, 3, 155–175.
19. Breyer, S. (1993). *Breaking the vicious circle: Toward effective risk regulation*. Cambridge, MA: Harvard University Press.
20. Baron, J., Wu, Z., Brennan, D. J., Weeks C., & Ubel, P. A. (in press). Analog scale, ratio judgment and person trade-off as utility measures: Biases and their correction. *Journal of Behavioral Decision Making*.
21. Baron, J. (in press). Measuring value tradeoffs: Problems and some solutions. In E. U. Weber, J. Baron, & G. Loomes (Eds.), *Conflict and tradeoffs in decision making: Essays in honor of Jane Beattie*. New York: Cambridge University Press.
22. Baron, J. (1996). Do no harm. In D. M. Messick & A. E. Tenbrunsel (Eds.), *Codes of conduct: Behavioral research into business ethics* (pp. 197–213). New York: Russell Sage Foundation.
23. Baron, J. (1998). *Judgment misguided: Intuition and error in public decision making*. New York: Oxford University Press.