

Gamma Discounting

By MARTIN L. WEITZMAN*

By incorporating the probability distribution directly into the analysis, this paper proposes a new theoretical approach to resolving the perennial dilemma of being uncertain about what discount rate to use in cost-benefit analysis. A numerical example is constructed from the results of a survey based on the opinions of 2,160 economists. The main finding is that even if every individual believes in a constant discount rate, the wide spread of opinion on what it should be makes the effective social discount rate decline significantly over time. Implications and ramifications of this proposed "gamma-discounting" approach are discussed. (JEL H43)

The concept of a "discount rate" is central to economic analysis, as it allows effects occurring at different future times to be compared by converting each future dollar amount into equivalent present dollars. Because of this centrality, the choice of an appropriate discount rate is one of the most critical problems in all of economics. And yet, to be perfectly honest, a great deal of uncertainty beclouds this very issue.

The problem of the unsure discount rate has long bedeviled benefit-cost analysis but has acquired renewed relevance lately because economists are now being asked to analyze environmental projects or activities the effects of which will be spread out over hundreds of years, and the evaluation of which is therefore extremely sensitive to the discount rate being used. Prominent examples include: global climate change, radioactive waste disposal, loss of biodiversity, thinning of stratospheric ozone, groundwater pollution, minerals depletion, and many others.¹

The most critical single problem with discounting future benefits and costs is that no consensus now exists, or for that matter has ever existed, about what actual rate of interest to use.

Economic opinion is divided on a number of fundamental aspects, including what is the appropriate value of an uncertain future "marginal product of capital," what are the relevant efficiency distortions and how are they possibly magnified by public-sector projects, how should large long-term public investments be placed in the framework of the capital-asset-pricing model, how are we to view intergenerational transfers when future generations are not presently represented, how do we account for equity and other distributional effects in aggregating costs, benefits, and discount rates over individuals or across countries or over time—and so forth, and so on. While economists may differ in their projections of distant future real interest rates (which depend, after all, on the ultimately unpredictable rate of technological progress), this is just the tip of the iceberg. More fundamental disagreements below the surface concern basic interpretations of welfare economics, including the role of governments and even the ethical foundations of intergenerational discounting. All these, and many more, considerations are fundamentally matters of judgement or opinion, on which fully informed and fully rational individuals might be expected to differ.

To some extent the economics profession deals with this disturbing state of affairs by looking the other way. First of all, the profession maintains, outsiders tend to exaggerate the actual differences among economists—there being an implicit premise here that a properly appointed committee of "experts" might be able to thrash out their differences in favor of a consensus social discount rate. Secondly, every

* Department of Economics, Harvard University, Cambridge, MA 01238. I am grateful to Robert Dorfman, David Laibson, and Robert M. Solow for their valuable comments. This work was supported in part by a grant from the National Science Foundation.

¹ For a recent discussion with further references, see William D. Nordhaus (1994) or Paul R. Portney and John P. Weyant (1999).

good economist knows that the proper procedure is to perform a sensitivity analysis using several plausible discount-rate values. Then, the logic continues, hopefully the right policy or project will be relatively robust to the discount rate being used for the evaluations.

The real situation, as we all know in our heart of hearts, is very different. There does not now exist within the economics profession, nor has there ever existed, anything remotely resembling a consensus, even—or, perhaps one should say, especially—among the “experts” on this subject. (Actually, with very little exaggeration or cynicism, an “expert” here might be defined as an economist who knows the literature well enough to be able to justify *any* reasonable social discount rate by some internally consistent story.²) Furthermore, and also contrary to the party line, expected net present discounted values for long-term projects are notoriously hypersensitive to the interest rate being used to make the evaluation. In a great many practical analyses of public projects, perhaps a majority, the most significant uncertainty of all concerns the discount rate itself.

This paper sets forth a radically different approach to the issue of what interest rate to use for discounting the future, starting with the recognition that the discount-rate problem is rooted in fundamental differences of opinion, which are unlikely to go away soon. No committee of “experts” will likely be able to resolve the dilemma—no matter how deeply they delve into it. Therefore, it follows, we should be operating from within a framework that incorporates the irreducible uncertainty about discount rates directly into our benefit-cost methodology.

The paper proposes an operational methodological framework to resolve the discount-rate dilemma, which is centered on the concept of “gamma discounting.” Viewed from within this framework, the more familiar exponential discounting is considered to be a special case of gamma discounting in the same spirit that an exponential probability distribution can be seen as a special case of the more general gamma probability distribution. (An exponential proba-

bility distribution is a one-parameter family describing the waiting time until breakdown of a single-component system, while the more general gamma probability distribution can be viewed as a two-parameter family representing the waiting time until breakdown of a multiple-component system, where the second parameter specifies the number of components in the system.)

The empirical part of this paper attempts to determine, from a survey of the opinions of 2,160 economists, the two defining parameters of the appropriate gamma distribution. It turns out empirically that the second, nonexponential, parameter of the gamma distribution plays, or at least *should* play, an extremely significant role in actual long-term discounting.

A numerical example is constructed, which suggests the practical desirability of thinking in terms of a future subdivided into approximately five subperiods. Named in order, these are, roughly: the *Immediate* Future (1 to 5 years hence), the *Near* Future (6 to 25 years hence), the *Medium* Future (26 to 75 years hence), the *Distant* Future (76 to 300 years hence), and the *Far-Distant* Future (more than 300 years hence). The numerical example suggests using the following approximation of within-period marginal discount rates for long-term public projects: Immediate Future about 4 percent per annum; Near Future about 3 percent; Medium Future about 2 percent; Distant Future about 1 percent; and Far-Distant Future about 0 percent.

The paper concludes by discussing some theoretical ramifications and practical implications of gamma discounting for global warming.

I. The Setting for the Problem

I begin this section by trying to avoid the bias that comes from being too close to see the forest for the trees. Let us stand back here to view the landscape as seen through the eyes of a hypothetical decision maker, here called the “WNO” (standing for “Wise Neutral Observer” or “World Nations Official”). The WNO is responsible for evaluating various projects or proposals to mitigate the possible effects of global climate change.

The WNO recognizes at once that a fundamentalist, full-blown, fully disaggregated, general-equilibrium-style analysis of the impact

² By *this* criterion, after conducting the survey of this paper I now rate myself as perhaps the world’s leading “expert” on the social discount rate.

of environmental projects is undoable, essentially because of analytical intractability and informational overload. Instead, the WNO follows most economists in opting for a practical compromise, which might be called the “pragmatic-decomposition” approach to intertemporal evaluation. The WNO decides to “pragmatically decompose” the evaluation of various projects by moving them through two specially appointed panels of expert appraisers.

Panel No. 1 makes period-by-period estimates of the aggregate net value of relevant income flow changes induced by a project, in terms of “expected consumption-equivalent” net real dollars for each period, and then delivers these estimates to the WNO. Such a task amounts to a truly formidable assignment, which is worthy of a more complete treatment. Yet this paper will not elaborate further the work of Panel No. 1—except to describe how each panel member’s individual opinion is aggregated into an overall expected-net-benefit estimate.

Panel No. 1 consists of m experts on evaluating global warming. Expert i ($1 \leq i \leq m$) consults model i , where a “model” is understood here in the broadest possible generic sense of “some mechanism” for predicting outcomes of projects to mitigate the possible effects of global climate change. Suppose that expert i (or model i) predicts a prototype project will yield expected consumption-equivalent net benefit (i.e., expected consumption-equivalent gross benefits minus costs) at time t of:

$$(1) \quad Z_i(t).$$

From the perspective of the WNO, each expert serving on Panel No. 1 is viewed as being of approximately equal standing in terms of objective credentials or any other observable marker of professional legitimacy. Therefore, implicitly using a diffuse uniform prior, what the WNO takes as the best statistical estimate of the appropriate *overall expected* consumption-equivalent net benefit at time t is simply the average response

$$(2) \quad Z(t) \equiv \frac{\sum_{i=1}^m Z_i(t)}{m}.$$

The $\{Z(t)\}$ from equation (2) constitutes simultaneously the fundamental *output* of Panel No. 1 and also the fundamental *input* of Panel No. 2. This paper also takes $\{Z(t)\}$ as given, and henceforth is focused almost exclusively on the work of the second panel.

Panel No. 2 is assigned the task of recommending the proper time-dependent weights $\{A(t)\}$ to be used for aggregating the net benefits from different times $\{Z(t)\}$ into a single overall index of *present-discounted expected consumption-equivalent net benefits*:

$$(3) \quad \int_0^{\infty} A(t)Z(t) dt.$$

Loosely speaking, then, the “principle of pragmatic decomposition” reduces the problem to evaluating a linear expression of the form (3), and then further subdivides the remaining task into two reasonably distinct subtasks. Panel No. 1 provides the time series of overall expected net benefits $\{Z(t)\}$. Panel No. 2 provides the corresponding time-evaluation aggregation weights $\{A(t)\}$. The WNO realizes full well that this procedure yields at best an approximate answer for what, in its full generality, constitutes an extraordinarily complex web of interrelated questions.³ Nevertheless, the “principle of pragmatic decomposition” seems like a satisfactory practical compromise to the WNO—especially when compared with the leading alternative, which is basically to curse the darkness by bemoaning the analytical and informational intractability of the problem. The WNO opts instead for lighting a small candle to illuminate a few modest guidelines for rough answers to some important questions.

Panel No. 2 consists of n experts on “the discount-rate question.” Expert j , ($1 \leq j \leq n$), consults model j , where here a “model” is understood in the broadest possible generic sense of “some mechanism” for generating time-

³ At best, the combining of $A(t)$ and $Z(t)$ in the linear form (3) can be rigorously justified only when the underlying $\{Z_i(t)\}$ are small relative to the world economy and independently distributed relative to the underlying $\{A_j(t)\}$.

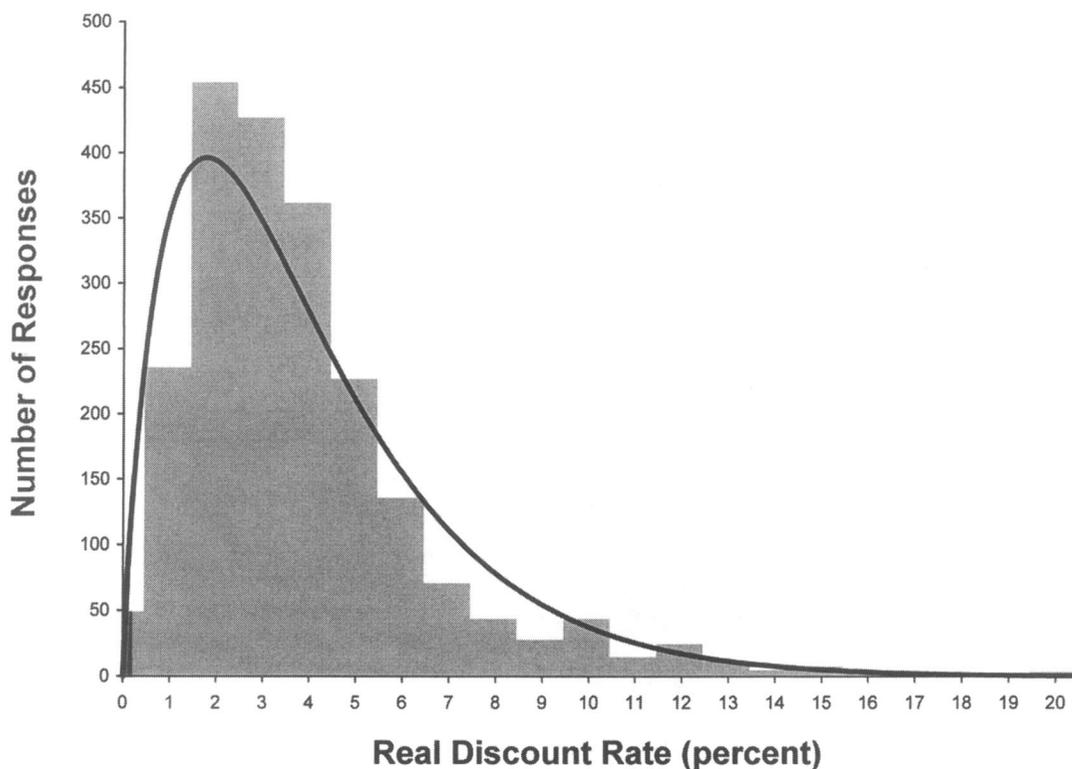


FIGURE 1. ACTUAL (HISTOGRAM) AND FITTED (CURVE) FREQUENCY DISTRIBUTIONS

evaluation weights. Suppose that expert j , or model j , recommends the time-evaluation aggregation weights:

$$(4) \quad A_j(t).$$

Like the first panel, then, members of this second panel also hold differing opinions—in this case about the discount-rate question. And, for this second panel also, the WNO has a diffuse prior on the credibility of each member's opinion. Only, here it is far less clear to the WNO how to combine these different opinions into a reasonable overall discounting function.

So far as the WNO can comprehend, social discounting weights apparently represent *opinions* more than anything else. The social discount rate here seems not so much to be a thing that economists *observe*, even in principle, but it appears rather more to represent a subjective *judgement* which they *believe in* to varying degrees. At this point, the WNO wisely realizes

that further practical progress will not be made without specifying more precisely the exact question to be asked each of the n experts serving on the second panel, and without placing somewhat more structure on the distribution of their various answers.

II. The Model

Imagine that a large number of economists have been asked to name the real-interest-rate value they favor for discounting long-term environmental projects. Suppose the replies are distributed approximately as the histogram shown in Figure 1. What might such an exercise suggest?

First of all, the large number of responses might suggest that the question is not altogether unfamiliar. It seems as if many economists are accustomed to thinking in terms of exponential discounting, but they exhibit a fair amount of variation in the preferred *value* of the constant discount rate that an individual economist is

more or less willing to recommend. To the WNO, desperate to render the problem more manageable, all of this suggests postulating the following enormously simplifying assumption.

ASSUMPTION 1 [Reduced-Form Specification #1]: *The individual time-evaluation aggregation weights $\{A_j(t)\}$ take the traditional exponential form*

$$(5) \quad A_j(t) = e^{-x_j t}.$$

Secondly, the histogram pattern suggests the general shape of a gamma probability distribution, like the smooth curve superimposed in Figure 1. Ever the empirical opportunist, the WNO accordingly makes the following assumption.

ASSUMPTION 2 [Reduced-Form Specification #2]: *The $\{x_j\}$ in (5) are distributed as the realizations of a random variable x , whose probability density function $f(x)$, defined for all positive x , is of the gamma form*

$$(6) \quad f(x) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\beta x}$$

with positive parameters α and β to be estimated from the data.⁴

In words, reduced-form specification #1 means that the average panel member knows about, and typically does not feel acutely uncomfortable with, the approximation of constant exponential discounting.⁵ The primary disagreement among panel members is over the appropriate value of the as-if-constant discount rate.

Reduced-form specification #2 then restricts the empirically observed discount-rate values to be distributed "as if" being realizations of the two-parameter gamma family. Equation (6) represents an openly opportunistic specification,

⁴ I assume some familiarity with the gamma distribution, which is typically described in any comprehensive book on probability theory. See, e.g., Morris H. DeGroot (1970).

⁵ To the extent that some panel members may believe in declining discount rates, the basic conclusions of this paper will only be strengthened. In this spirit, the main conclusion might be stated as follows: *Even if everyone believes in a constant discount rate, the effective discount rate declines strongly over time.*

because it reduces the WNO's problem to pinning down the parameter values of α and β . Imposing the gamma distribution here is a strategic assumption that will deliver enormous analytical tractability at very little cost in terms of actual empirical restrictiveness.

What is the expected value today of an extra expected dollar at time t ? It should be the expected present discounted value of a dollar at time t , weighted by the "probability of correctness" or the "probability of actuality" of the rate at which it is being discounted:

$$(7) \quad A(t) \equiv \int_0^\infty e^{-xt} f(x) dx.$$

The function $A(t)$ may be called the *effective discount function for time t* . Note that the WNO is essentially taking a probability-weighted average (over all types of persons in all states of mind) of discount functions of the members of Panel No. 2. The key insight here is that what should be probability-averaged at various times is not discount rates, but discount functions.

The corresponding marginal or instantaneous effective discount rate at time t is defined to be

$$(8) \quad R(t) \equiv -\frac{\dot{A}(t)}{A(t)}.$$

For the gamma probability distribution (6), the effective discount function defined by (7) is

$$(9) \quad A(t) = \frac{\beta^\alpha}{\Gamma(\alpha)} \int_0^\infty x^{\alpha-1} e^{-(\beta+t)x} dx.$$

Now it is a fact of calculus that

$$(10) \quad \int_0^\infty y^{a-1} e^{-by} dy = \frac{\Gamma(a)}{b^a}.$$

By substituting the values $a = \alpha$ and $b = \beta + t$ into formula (10), and then using the resulting expression, equation (9) can be transformed into

$$(11) \quad A(t) = \left(\frac{\beta}{\beta + t} \right)^\alpha.$$

Now the mean and variance of (6) are defined in the usual way

$$(12) \quad \mu \equiv \int_0^\infty x f(x) dx$$

and

$$(13) \quad \sigma^2 \equiv \int_0^\infty (x - \mu)^2 f(x) dx.$$

Making use again of (10), it can be proved for the gamma distribution that

$$(14) \quad \mu = \frac{\alpha}{\beta},$$

and

$$(15) \quad \sigma^2 = \frac{\alpha}{\beta^2}.$$

Because μ and σ have economically intuitive meanings, while α and β have no economic significance per se, it will be more appropriate in this paper to work with the inverse of (14) and (15), namely

$$(16) \quad \alpha = \frac{\mu^2}{\sigma^2}$$

and

$$(17) \quad \beta = \frac{\mu}{\sigma^2}.$$

After substituting from (16) and (17), equation (11) can be rewritten as

$$(18) \quad A(t) = \frac{1}{(1 + t\sigma^2/\mu)^{\mu^2/\sigma^2}}.$$

Equation (18) can readily be translated into an equivalent formulation that makes clear the exact sense in which gamma discounting may be viewed as a generalization of exponential discounting. Let N be a fixed parameter standing for the number of discrete times at which interest payments are to be compounded within any time interval. If μ represents the relevant interest flow rate, then, with N equally spaced compounding operations in the interval $[0, t]$, a dollar on deposit at time zero grows into

$$(19) \quad B(t) = \left(1 + \frac{\mu t}{N} \right)^N$$

dollars at time t .

But now notice, comparing (18) with (19), that

$$(20) \quad A(t) = \frac{1}{B(t)},$$

whenever we interpret N to be defined by

$$(21) \quad N \equiv \left(\frac{\mu}{\sigma} \right)^2.$$

Exponential discounting can thus be viewed as a special case of gamma discounting appropriate to a limiting situation where the number of compounding operations N is very large, which by (21) can be interpreted to be essentially the same thing as having σ be very small.

As economists, we are perhaps unaccustomed to thinking in terms of a situation where N is held constant in the familiar equation (19), while t is varied. But there is nothing inherently amiss in the image of a world acting like a bank that limits the "interest paid on interest" by restricting itself to just N compounding operations performed within any time interval $[0, t]$. Such a "gamma bank" goes by the rule that if you deposit a dollar at time zero and then close out your account at time t , the withdrawal value of your savings will be compounded exactly N times in the interval $[0, t]$, irrespective of the withdrawal time t . In this interpretation the amount of each interest payment $\mu t/N$ credited per dollar on deposit varies linearly with

withdrawal time t , while the total *number* of such interest compounding operations N is viewed as being fixed beforehand.

Returning to the primary formulation (18), by equation (8) the implied effective discount rate is

$$(22) \quad R(t) = \frac{\mu}{1 + t\sigma^2/\mu}.$$

It was already noted that exponential discounting corresponds here to the special limiting case of (18) and (22) when the parameter σ approaches zero. Note also from (18) that, for any fixed positive value of σ , the effective discount rate starts off at its mean value μ when $t = 0$, but declines monotonically towards zero over time.⁶ The effective discount rate declines with the passage of time because increasingly greater relative present-value weight is being placed on the states with the lower rates. In comparison, the higher-discount-rate states are relatively less important over time, because their present value has been more drastically shrunk by the power of compound discounting at the higher rates.

What are the policy implications of time-varying discount rates? It depends, ultimately, on what is being assumed about how knowledge of future interest rates is revealed sequentially. If the expert panel remains equally divided as time unfolds, because each resolution of interest-rate uncertainty is balanced by a new source of uncertainty, there is a potential problem of time inconsistency because, for sequential decision-making problems, nonexponential discounting may lead to investment plans that do not fully cohere over time.⁷

I think it is fair to characterize this paper as suggesting empirically that long-run time inconsistency may run deeper and be more generic than was previously suspected. However, the implications of time inconsistency per se are properly the subject of another paper. In the

context of this paper, I think it suffices to think of formula (18) [or (19)–(21)] as defining a table of technocratic time-dependent weights, which give future-dollar equivalence values for making one-time irreversible investment decisions. The possibility of being able later to revisit and revise time-inconsistent investment choices introduces a set of complicated issues that are better treated separately.

Equations (18) and (22) represent the theoretical or conceptual centerpiece of the paper, and provide a starting point for the empirical estimation of $\{A(t)\}$ and $\{R(t)\}$. In the form of (18) and (22) we have obtained a simple, useful, two-parameter closed-form expression that generalizes exponential discounting to cover a situation of uncertain discount rates.

III. The “Professionally Considered Gut Feeling” of Economists

We next attempt to push the analysis one step further by deriving actual numerical estimates for the declining effective discount-rate schedules appropriate to evaluating long-term environmental projects, like mitigating the effects of global climate change. Building on equations (18) and (22), what essentially remains is to fix empirically the two parameters μ and σ .

This section proceeds as if two surveys were conducted. The first survey represents a large-scale unscreened sampling of the opinions of over 2,000 professional Ph.D.-level economists. The second survey effectively comprises a named subsample of 50 representative blue-ribbon “leading economists.”

The full text is reproduced in the Appendix, but the operative part of the questionnaire asked the subjects to reply with their “*professionally considered gut feeling*” to the following request:

Taking all relevant considerations into account, what real interest rate do you think should be used to discount over time the (expected) benefits and (expected) costs of projects being proposed to mitigate the possible effects of global climate change?

The questionnaire was sent by e-mail to about 2,800 Ph.D.-level economists. Efforts were made to ensure that the sample represented a balance of backgrounds, fields, and viewpoints.

⁶ This result justifying hyperbolic discounting is not restricted to the gamma form, but holds generally for any probability distribution. See Weitzman (1998) for more details.

⁷ For a further exposition of this problem, see, e.g., Maureen L. Cropper and David Laibson (1999).

Additionally, I tried to include all those whom I consider to be the very best economists around, irrespective of background or area. Over 2,100 usable replies were received from economists in some 48 countries, representing every major field of economics.

I hasten to add that about one in eight respondents objected to some aspect or another of the question they were being asked. What respondents did *not* like varied considerably. Some wanted to be able to give two different interest rates—one for developed and a second for underdeveloped countries. For others, the lumping together of developed and underdeveloped countries was all right, but they wanted different rates to be applied separately to costs and benefits, or to different projects, or to different parts of different projects, or to something else.

Another set of objections centered on the need to take explicit account of distributional considerations, with some of these respondents seemingly wanting me to supply them with an appropriate set of income-distribution welfare weights. Still another group of disgruntled participants wanted, in effect, to be supplied with the appropriate “beta” coefficient for the particular investment project in mind, since, according to them, it is impossible to speak meaningfully about a general discount rate “for projects to mitigate the effects of global climate change” without specifying more explicitly how the payoffs are supposed to be correlated with the performance of some index of alternative investments, like the value of the world’s stock markets.

Many of these objections struck me as being legitimate, at least in principle. The various types of complaints actually tended to have relatively little in common with each other, aside from the shared bottom-line feeling that the question as it stood was problematical, or perhaps even meaningless. It seems as if what sometimes bothered one economist to an expression of indignation about the survey typically failed to arouse another at all. For almost every case of a raised objection, I acted like the sitting judge who wants badly to avoid a hung jury—I tried to make these sophisticated-sounding jurist-respondents give an actual numerical answer for the imperfect world where we happen to live, which is dominated by simpleminded people, like me, who can only un-

derstand simplistic answers to complicated questions.

Overall, the most common objection from respondents was their complaint that they were “not an expert in this area” and consequently “have no idea” what to answer. (The most extreme position of this sort held that economists should not be asked to express an opinion at all; rather it should be politicians or somebody else better qualified to reflect public opinion.) Once again, like the tunnel-visioned judge seeking to avoid a hung jury, I assured respondents of this type that there were no real experts out there, only folks like themselves with honest differences of opinion—some of whom just happen to be more articulate than others in crafting stories to justify their answers. I would then typically try to prevail upon these concerned “non-experts” to give their “professionally considered gut feeling” for the imperfect world in which we live, effectively asking them as Ph.D.-level economists to integrate out of the process their subjective uncertainties concerning whatever they “have no idea” about.⁸

For “extreme” responses of less than 0.5 percent or greater than 12 percent, I typically requested confirmation and asked for some brief rationalization. If the respondent confirmed with *any* remotely plausible story (and held a Ph.D. degree in economics from a recognizable graduate school), the response was recorded.

At 2,160 replies from economists in 48 countries, representing all major fields of economics, the coverage was comprehensive and the overall response rate was excellent. But, the reader should be warned again, approximately 12

⁸ For those very few respondents who volunteered to report a range of discount-rate values, over which parametric analysis could be performed, I typically reacted by acknowledging that they were, of course, correct in principle, but I asked them nevertheless to pinpoint a single number as if a simpleminded policy maker could only use one value. Had I thought in the first place that a large number of respondents might cooperate analogously, by revealing a distribution of their beliefs, I could have designed the survey question along such lines from the beginning; however, as was confirmed by the difficulties I subsequently experienced in getting some economists to comply with revealing just a single number, I feared that requests for any more information, even for a range of values, might have met with refusals to cooperate at all, which could have introduced a significant element of response bias.

TABLE 1—DISTRIBUTION OF RESPONSES

Discount rate (Rounded to nearest whole percentage)	Number of responses
-3	1
-2	1
-1	1
0	46
1	236
2	454
3	427
4	362
5	227
6	136
7	71
8	44
9	28
10	44
11	15
12	25
13	12
14	5
15	8
16	3
17	2
18	3
19	1
20	4
25	2
26	1
27	1
Total responses = 2,160	

percent of the replies might legitimately be classified in the category of “answering under duress” or “responding under objection.” My opinion was, and remains, that it is far better to have a sample relatively free of response bias than to rely on self-styled “experts” who, in my experience, are sometimes bluffing and, in any case, may have all kinds of agendas hidden behind their willingness to thrust forward “their” number. Accordingly, I invariably prodded unsure respondents, sometimes to the point of hectoring, to come up with some “best estimate” rather than to leave the profession with the results of a survey based upon a biased sample.

The sample mean is $\mu = 3.96$ percent, with a standard deviation $\sigma = 2.94$ percent. In Table 1 all 2,160 responses are listed, rounded off to the nearest whole integer, and the frequency distribution is depicted as the histogram of Figure 1. The corresponding gamma probability distribution, which is defined by the same pa-

rameter values of μ and σ as the sample, is the smooth curve drawn in Figure 1. (For visual comparison, both distributions are scaled to have the same area, namely $n = 2,160$.) Even without digressing into fancy nonlinear statistics by attempting here to define formally “goodness of fit,” I think that the resemblance in Figure 1 between the empirical histogram and the theoretical curve is, intuitively, a “sufficiently reasonable match” to allow, for the practical purposes of this paper, the very considerable analytical convenience of the gamma form. Thus, it may be noted in passing, the statistical methodology of this study is simple in the extreme.

Any one particular economist may feel somewhat unsure about what answer to give, but the enormous size of the sample allows there to be revealed a great deal of underlying statistical regularity in the aggregate responses. We are looking in Table 1 at a highly skewed distribution, showing a lot of variation, with a mean of about 4 percent, a median of 3 percent, and whose mode is 2 percent. Note the pronounced “round-off effect” for “familiar” large integer rates, like 10 percent, 12 percent, 15 percent, or 20 percent.

I also conducted an “as-if” second survey, which is just a broken-out subsample of 50 high-profile named economists, selected to constitute a hypothetical “balanced blue-ribbon panel” of expert opinion. In alphabetical order, the 50 leading economists chosen to constitute the “as-if” expert panel are:

George AKERLOF; Kenneth ARROW; Anthony ATKINSON; Robert BARRO; William BAUMOL; Gary BECKER; David BRADFORD; John CAMPBELL; David CARD; William CLINE; Peter DIAMOND; Avinash DIXIT; Jacques DREZE; Martin FELDSTEIN; Stanley FISCHER; Roger GORDON; Robert HALL; Arnold HARBERGER; Jerry HAUSMAN; Dale JORGENSON; Lawrence KOTLIKOFF; David KREPS; Paul KRUGMAN; Jean-J. LAFFONT; Hayne LELAND; Robert LUCAS; Karl-G. MALER; Burton MALKIEL; Alan MANNE; Daniel McFADDEN; Robert MERTON; James MIRRLEES; Kevin MURPHY; William NORDHAUS; Pierre PESTIEAU; Paul PORTNEY; James POTERBA; Stephen ROSS; Agnar SANDMO; Thomas

SCHELLING; Richard SCHMALENSEE; Myron SCHOLES; Amartya SEN; Robert SHILLER; John SHOVEN; Robert SOLOW; Joseph STIGLITZ; Lawrence SUMMERS; James TOBIN; W. Kip VISCUSI.

I cannot see how it might be said fairly that the 50 members of this group are not distinguished economists or are ignorant of the issues involved—any more than it could be argued that the full sample of 2,160 names is not sufficiently comprehensive to uncover the underlying probability distribution.

The aggregate response for the above-listed “balanced blue-ribbon panel” of expert opinion is: $\mu = 4.09$ percent; $\sigma = 3.07$ percent. Because the mean and standard deviation for the entire sample of 2,160 professional economists coincides so closely with the mean and standard deviation for the subsample of 50 blue-ribbon “leading-expert” economists, I take advantage of the opportunity here to average the two results and round off to the nearest one-tenth of one percent. Henceforth in this paper, I will work with the numerical values:

(23) $\mu = 4$ percent per annum

and

(24) $\sigma = 3$ percent per annum.

IV. Gamma Discounting for Global Warming

From equation (22), the schedule of effective discount rates is a simple rectangular hyperbola that is easy to calculate. What might be called its *elasticity of discount-rate decline* is

(25)
$$-\frac{t\dot{R}(t)}{R(t)} = \frac{1}{1 + \frac{1}{\nu t}}$$

where the coefficient

(26)
$$\nu \equiv \frac{\sigma^2}{\mu}$$

can be interpreted as a “convergence velocity” parameter. One can then speak of the effective

discount-rate schedule declining hyperbolically to its asymptote of zero at a “convergence velocity” given by formula (26).

It is important to understand why the effective discount rate declines more rapidly when the underlying spread of professional opinion is broader. For the same mean, a greater variance of opinion signifies that more professionals believe in the *lower*-rate scenario, which is the scenario that predominates over time in present-value calculations. Of course, more professionals then also believe in the *higher*-rate scenario, but *their* present-value weight becomes relatively less significant over time, due to the power of compound discounting at the higher rate. (In a sense, the high-rate believers discount away the relevance of their own scenario, leaving the future ultimately to the low-rate believers.) Summing the two effects, a greater spread of opinion hastens the day when the low-rate scenario predominates in present-value calculations, thus causing the effective discount rate to decline more rapidly.⁹

Equation (22) can readily be inverted to yield the expression

(27)
$$t(R) = \frac{\frac{\mu}{R} - 1}{\nu}$$

where the velocity parameter ν is defined by (26).

With the numerical values implied by (23), (24), and (26), formula (27) can be used to determine the time cutoff values corresponding to $R = 3.5$ percent, $R = 2.5$ percent, $R = 1.5$ percent, and $R = 0.5$ percent. After rounding off, this numerical exercise then suggests a very rough periodization of the future along the approximate guidelines expressed in Table 2.

Table 2 speaks for itself in presenting the principal policy implication of this paper. The existing discount-rate uncertainty generates a sliding-scale effective discount-rate schedule, whose decline over time is significant enough

⁹ I am unsure how to characterize the most general class of probability distributions for which higher σ implies lower $R(t)$. I conjecture the condition holds for the class of all distributions for which σ parameterizes a mean-preserving spread, but at this stage cannot provide a rigorous proof.

TABLE 2—"APPROXIMATE RECOMMENDED" SLIDING-SCALE DISCOUNT RATES

Time period	Name	Marginal discount rate (Percent)
Within years 1 to 5 hence	<i>Immediate Future</i>	4
Within years 6 to 25 hence	<i>Near Future</i>	3
Within years 26 to 75 hence	<i>Medium Future</i>	2
Within years 76 to 300 hence	<i>Distant Future</i>	1
Within years more than 300 hence	<i>Far-Distant Future</i>	0

that this feature ought to be incorporated into any discounting of long-term environmental projects. With numerical values (23) and (24), the velocity parameter (26) is sufficiently large to mandate the explicit use of a time-dependent discount formula, such as (18), (22), for many practical applications of benefit-cost analysis.

The downward-sloping time profile of the discount-rate schedule displayed in Table 2 is sufficiently steep to make the equivalent "as-if-constant" discount rate much lower than the average discount rate of 4 percent. To see this effect clearly in a specific example, suppose that an investment project is of the familiar point-input continuous-output form. Then the equivalent "as-if-constant" discount rate is

$$(28) \quad \bar{r} = \frac{1}{\int_0^{\infty} A(t) dt}.$$

Substituting for $A(t)$ from (18) and carrying out the integration in (28) yields

$$(29) \quad \bar{r} = \frac{(\mu - \sigma)(\mu + \sigma)}{\mu}.$$

With the values $\mu = 4$ percent and $\sigma = 3$ percent being assumed in (23) and (24), expression (29) becomes

$$(30) \quad \bar{r} = 1.75 \text{ percent per annum.}$$

The numerical value (30) therefore implies the following. If we were forced here to choose a single constant-equivalent discount rate \bar{r} to represent the entire sliding-scale schedule for a

point-input continuous-output investment, then it would be less than 2 percent per annum.

What is this paper then implying about the optimal *form* of a long-term project under uncertainty (such as taking measures to ameliorate global warming)? The full answer is complicated and conditional on the exact time profile of the investment, because of the downward tilt of the effective discount rate as a function of time. But the decline is sufficiently steep to suggest, at some considerable risk of oversimplification, the following sort of generalization. Other things being equal, the basic proposition should make itself felt by biasing the choice of policy instruments (such as taxes vs. tradeable permits) and levels of desired stringency (such as greenhouse gas emission targets) *as if* toward what is optimal for a low-interest-rate scenario because, other things being equal, that kind of scenario will weigh more heavily in the expected present-discounted difference between benefits and costs.¹⁰

V. Conclusion

The very wide spread of professional opinion on discount rates means that society should be using effective discount rates that decline from a mean value of, say, around 4 percent per annum for the immediate future down to around zero for the far-distant future. Furthermore, the decline in effective social discount rates is sufficiently pronounced, and comes on line early

¹⁰ This seems like a roughly accurate description of some of the main findings in William A. Pizer (1999). Please note I am *not* asserting here that everything can be reduced to an equivalent constant interest rate. It is only that *if* we force ourselves to think in such oversimplified terms, which is wrong to begin with, *then* the best compromise is to think in terms of a low rate.

enough, to warrant inclusion of this sliding-scale feature in any serious benefit-cost analysis of long-term environmental projects, like activities to mitigate the effects of global climate change.

APPENDIX: TEXT OF QUESTIONNAIRE

Dear XXX:

I would very much appreciate enlisting a few moments of your time in the cause of introducing some economic analysis into the current policy debates about global warming.

As part of an empirical study, I need your best point estimate of the appropriate real discount rate to be used for evaluating environmental projects over a long time horizon. (What I am after here is the relevant interest rate for discounting real-dollar changes in future goods and services—as opposed to the rate of pure time preference on utility.)

Your response will be held in strictest confidence. Only highly aggregated statistics will be reported. No one else will find out what number you gave me.

Try to imagine, then, that an international organization, such as the U.N. or the World Bank, has commissioned some well-done, comprehensive, and highly-professional studies on the possible quantitative impacts of global climate change. Suppose these studies have done a credible job of converting costs and benefits from any given year into expected-consumption-equivalent real dollars for that year. Now these organizations want your advice—and are possibly going to take it seriously—about what discount rate to use for aggregating together the expected-consumption-equivalent real dollars from different years, in order to calculate the overall net present discounted value of a proposed project.

You are being asked the following question. *“Taking all relevant considerations into account, what real interest rate do you think should be used to discount over time the (expected) benefits and (expected) costs of projects being proposed to mitigate the possible effects of global climate change?”*

Your answer should be a single number. (I realize that a single constant discount rate here

represents a vast oversimplification, but you can pretend that this is the only language that policy makers understand.)

Naturally, I would prefer to have your best answer be a seriously-considered and well-thought-out estimate. However, an even higher priority in the study must be placed on attaining extremely high response rates, which are required here for statistical accuracy by the underlying experimental design. Therefore, in order to eliminate response bias from the sample sub-group that you represent, I will settle for what might be called your “professionally considered gut feeling.”

I know that you probably don’t consider yourself to be an expert in this area, but I want your best opinion anyway. Please respond to this request by return e-mail, even if your preferred number is a back-of-the-envelope guesstimate—or even if it just comes off the top of your head. No explanation is required.

Thank you for your help.

Marty Weitzman

mweitzman@harvard.edu

REFERENCES

- Cropper, Maureen L. and Laibson, David.** “The Implications of Hyperbolic Discounting for Project Evaluation,” in Paul R. Portney and John P. Weyant, eds., *Discounting and intergenerational equity*. Washington, DC: Resources for the Future, 1999, pp. 163–72.
- DeGroot, Morris H.** *Optimal statistical decisions*. New York: McGraw-Hill, 1970.
- Nordhaus, William D.** *Managing the global commons: The economics of climate change*. Cambridge, MA: MIT Press, 1994.
- Pizer, William A.** “The Optimal Choice of Climate Change Policy in the Presence of Uncertainty.” *Resource and Energy Economics*, August 1999, 21(3–4), pp. 255–87.
- Portney, Paul R. and Weyant, John P., eds.** *Discounting and intergenerational equity*. Washington, DC: Resources for the Future, 1999.
- Weitzman, Martin L.** “Why the Far-Distant Future Should Be Discounted at Its Lowest Possible Rate.” *Journal of Environmental Economics and Management*, November 1998, 36(3), pp. 201–08.