Journal of Economic Literature
Vol. XLV (September 2007), pp. 703–724

A Review of *The Stern Review on the Economics of Climate Change*

Martin L. Weitzman*

The Stern Review *calls for immediate decisive action to stabilize greenhouse gases because "the benefits of strong, early action on climate change outweighs the costs." The economic analysis supporting this conclusion consists mostly of two basic strands. The first strand is a formal aggregative model that relies for its conclusions primarily upon imposing a very low discount rate. Concerning this discount-rate aspect, I am skeptical of the* Review's *formal analysis, but this essay points out that we are actually a lot less sure about what interest rate should be used for discounting climate change than is commonly acknowledged. The* Review's *second basic strand is a more intuitive argument that it might be very important to avoid possibly large uncertainties that are difficult to quantify. Concerning this uncertainty aspect, I argue that it might be recast into sound analytical reasoning that might justify some of the* Review's *conclusions. The basic issue here is that spending money to slow global warming should perhaps not be conceptualized primarily as being about consumption smoothing as much as being about how much insurance to buy to offset the small change of a ruinous catastrophe that is difficult to compensate by ordinary savings.*

1. *Introduction*

The issue of global climate change and what to do about it has put economics to a severe test in which economists have been challenged to think afresh about how to model (or at least how to conceptualize) such fundamental notions as risk, uncertainty, and discounting. There is nothing like being asked for a specific policy recommendation on a vivid actuality to breathe new life into otherwise arcane matters of economic analysis. Beyond the issue of whether it is right or wrong in its conclusions, the *Stern Review on the Economics of Climate Change* is an opportunity for economists to take stock of what we know about this subject, how we know it, what we don't know, and why we don't know it.

The *Stern Review* is a full-fledged economic analysis of climate change that was officially commissioned by the British government and, for reasons both economic and political, is an unusual—and unusually important—document. Sir Nicholas Stern is a professional economist of high standing and a distinguished public servant. Weighing in at close to 700 pages, the *Stern Review* is

* Weitzman: Harvard University. For helpful detailed comments on earlier drafts of this paper, but without implicating them for its remaining defects, I am grateful to Scott Barrett, Roland Benabou, Olivier Blanchard, Richard Cooper, Stephen DeCanio, Howard Gruenspecht, Cameron Hepburn, Chris Hope, Donald Ludwig, Robert Mendelsohn, Larry Samuelson, Robert Solow, Nicholas Stern, Lawrence Summers, and Hal Varian.

comprehensive in its scope and ambitious in its aims, with an attractive multicolored visual design that makes topics like cost–benefit analysis of dynamic externalities look almost glamorous. Anyone wanting to get a good feel for the basic issues of global climate change could profitably browse through this report, which covers well its multiple facets in a reader-friendly format. The *Review* contains much of value and interest aside from its cost–benefit analysis of mitigation policies, although that is naturally the part which most grabs the attention of economists. A detailed *Review of the Review* is out of place here—it would be too long, and besides the *Stern Review* reads well and is available online. Instead, I concentrate here on trying to distill the *Review* down to what I think is its analytical essence as a piece of applied cost–benefit analysis, because there can be difficulty seeing the forest for the trees when there are so many trees.

To make a long story short, the *Stern Review* comes down very strongly on the side of undertaking decisive—and expensive—measures starting now to reduce CO_2 and other greenhouse gas emissions because (and this quote captures well the tone of urgency about moving quickly to avoid catastrophic possibilities that is evident throughout the report): "Our actions over the coming few decades could create risks of major disruption to economic and social activity, later in this century and in the next, on a scale similar to those associated with the great wars and economic depression of the first half of the 20th century" (p. xv). Such a strong call to immediate decisive action is at odds with what most other economic analyses of climate change have concluded. The majority view of most other economic analysts finds it optimal to pursue a more gradualist course by starting with greenhouse gas emissions reductions at far lower levels than what the *Stern Review* advocates for the near future, but which after that ramp up considerably over time. The *Review* analysis, on the other hand, finds that "the benefits of strong and early action far outweigh the economic costs of not acting" (p. xv) and calls for stabilizing greenhouse gas atmospheric concentrations at ≈ 550 parts per million (ppm) of CO_2-equivalent (CO_2e). (The current level is ≈ 430 ppm CO_2e, compared with ≈ 280 ppm CO_2e before the Industrial Revolution.) This would make temperatures a hundred years from now be at $E[\Delta T] \approx 2°C$ and would (hopefully) stabilize future temperatures permanently thereafter at $\Delta T \approx 3°C$. By contrast, along the more gradual majoritarian optimal trajectories CO_2e concentrations a century from now are > 600 ppm and $E[\Delta T] \approx 2.5°C$—with temperatures expected to continue rising to well *above* $E[\Delta T] \approx 3°C$ *after* year 2105. To accomplish the *Review*'s ambitious goal, greenhouse gas emissions would need to be progressively cut by ≈ 3 percent each year, beginning more or less immediately. Which brings us to a central question. Why is there such a big difference between what *Stern* is recommending and what most other serious analysts favor?

This paper makes five basic points about the economics of climate change: (1) the discount rate we choose is all important and *Stern*'s results come from choosing a very low discount rate; (2) we are a lot less sure about core elements of discounting for climate change than we commonly acknowledge because critical puzzles, projections, and ambiguities are yet unresolved; (3) standard approaches to climate change (even those that purport to treat uncertainty) fail to account fully for the implications of large consequences with small probabilities; (4) structural parameter uncertainty that manifests itself in the thick tails of reduced-form probability distributions—not risk—is what likely matters most; (5) gathering information about thick-tailed uncertainties representing rare climate disasters (and developing a realistic emergency plan were they to materialize) should be a priority of research. To anticipate my main finding, spending money now to slow global warming should not be conceptualized primarily as being about optimal consumption smoothing so much as an issue about how much insurance to buy to offset the small chance of a ruinous catastrophe that is difficult to compensate by ordinary savings. While I am (along with most other economist-critics) skeptical of *Stern*'s formal analysis, I believe that the *Review*'s informal emphasis on climate-change uncertainty can be recast into sound analytical arguments that might justify some of its conclusions.

2. *Interest Rates and Long-Term Discounting*

Overall, I believe it is fair to say that the *Stern Review* consistently leans toward (and consistently phrases issues in terms of) assumptions and formulations that emphasize optimistically low expected costs of mitigation and pessimistically high expected damages from greenhouse warming—relative to most other studies of the economics of climate change. But far more crucially, the key assumption that drives its strong conclusions is the mundane fact that a very low interest rate is postulated, with which distant-future benefits and costs are then discounted. The upward-sloping "climate policy ramp" of ever-tighter emissions reductions in the majority of other models (but not beginning just yet, please) is a familiar consumption-smoothing consequence of discounting: the higher the interest rate the stronger the desire to move toward getting more pleasure now at the expense of postponing more pain until later. An efficient trajectory has a cost minimizing substructure similar to a Hotelling extraction problem: consumption flows are smoothed over time by maximizing present discounted utility subject to a stock constraint on accumulated CO_2e, which results in an "as if" CO_2e shadow tax that grows over time at (approximately) the rate of interest. The *Stern Review* simultaneously raises overall greenhouse gas reductions and flattens the "climate policy ramp" in its Hotelling-analogous consumption smoothing time profile by imposing discounting at a bare-minimum rate of interest.

Global climate change unfolds over a time scale of centuries and, through the power of compound interest, what to do now is hugely sensitive to the discount rate that is postulated. In fact, it is not an exaggeration to say that the biggest uncertainty of all in the economics of climate change is the uncertainty about which interest rate to use for discounting. In one form or another, this little secret is known to insiders in the economics of climate change, but it needs to be more widely appreciated by economists at large. The insight that the strong conclusions of the *Review* are driven mainly by the low assumed discount rate has been picked up and commented upon already by several insider critics.[1] Here I want to paraphrase this important debate for outsider economists and in the process bring some new ingredients to the mix.

An *Integrated Assessment Model*—hereafter IAM—is insider lingo for a multiple-equation computer-simulated model that combines dynamic economics with geophysical climate dynamics for the purposes of analyzing the economic effects of global climate change.[2] An IAM is essentially a model of economic growth with a controllable externality of endogenous greenhouse warming. The *Review* uses an IAM called PAGE, on which some numbers have been crunched and some conclusions have been based, but the exact connection between PAGE and *Stern*'s conclusions is elusive, frustrating, and ultimately unsatisfactory for a professional economist who honestly wants to understand where the strong policy recommendations are coming from.[3] The analytical core of the

[1] Variants of this argument are made in Partha Dasgupta (2007), Robert O. Mendelsohn (2007), William D. Nordhaus (2007), and Richard S. J. Tol and Gary W. Yohe (2006).

[2] A survey of integrated asssessment models for climate change control is provided in David L. Kelly and Charles D. Kolstad (1999).

[3] A nice description of how the *Stern Review* uses (and misuses) PAGE is contained in David Maddison (2006). PAGE itself is described in Chris Hope (2006).

Review is chapter 6 ("Economic Modelling of Climate-Change Impacts"), which is loosely tied to PAGE. However, the rest of the book contains lots of stories and examples suggesting that difficult-to-quantify uncertainty about really bad climate extremes may actually be an important informal part of *Stern*'s overall case. Economists are justifiably suspicious when someone refuses to aggregate various probability-weighted scenarios into an overall cost–benefit assessment, which at least can serve as a conversation starter. (How else are we to evaluate overall policy advice, such as what *Stern* recommends to us, except in the context of some overall model where assumptions and specifications are spelled out clearly?) As economic analysis, the *Stern Review* dwells in a nonscientific state of limbo where it uses an IAM but simultaneously refuses to commit to it or to any other consistent overarching framework within which its radical recommendations might be deconstructed and judged by others. Instead, the *Review* dances around the significance of the aggregative analysis of chapter 6 by arguing that conclusions from IAMs are suggestively useful but not crucial to the basic story line that anything above ultimate stabilization at ≈550 ppm of CO_2e and $\Delta T \approx 3°C$ is self-evidently just too risky for the planet to bear. In trying to make some overall sense of *Stern*'s mixed methodology (called "multidimensional" in the *Review*), I propose here to lay out the core issues of how risk, uncertainty, and discounting interact with the economics of climate change in terms of the simplest general-equilibrium model I can think of. Then I will try to clothe some parts of *Stern*'s intuitions about climate-change uncertainty in formal garb.

Irving Fisher taught us that an interest rate, like any other price, is the outcome of a dynamic general-equilibrium interaction of tastes with technology. The modern incarnation of Fisher's idea in a deterministic setting is the famous Frank Ramsey equation

$$r = \delta + \eta g, \tag{1}$$

where r is *the* interest rate (more on *the* interest rate later), δ is the rate of pure time preference, g is the per capita growth rate of consumption, and η is the elasticity of marginal utility, or, equivalently, the coefficient of relative risk aversion. In the shorthand notation of (1), the parameters δ and η capture two critical aspects of "tastes" (or "preferences") while the reduced-form representation of "technology" is the growth rate of consumption g. The important distinction between δ and r is that δ is a more primitive rate of pure time preference that *discounts utility*, while r is the much more familiar interest rate used to *discount consumption*, which is derived from *all* of the more primitive underlying parameters of tastes and technology via (1). The other taste parameter η represents the relative curvature of the utility function and is simultaneously a measure of aversion to interpersonal inequality and a measure of personal risk aversion. On the technology side, formula (1) holds whether g is endogenous or exogenous. In Ramsey's time, g was conceptualized as coming from capital accumulation, and therefore in long run equilibrium with diminishing returns to capital $g \to 0$ and $r \to \delta$. We now know from modern post-Solow growth theory (but Ramsey and Fisher didn't) that, in balanced growth steady-state equilibrium, g is essentially the underlying growth rate of labor-augmenting technological progress that, behind the scene, is pushing the entire economy forward (at least in a world without a greenhouse-warming externality). What I propose to do here is use the Ramsey equation as a transparency-based springboard for recasting the economics of climate change in terms of the four critical variables that appear in (1): δ, η, g, and r. I will ultimately argue that, in a greenhouse gas world, g needs to be seen as a random variable whose probability distribution has a climate-change-thickened left tail that carries most of the weight of expected marginal utility in cost–benefit analysis.

To cut sharply to the essence of the core discounting issue behind the *Review*'s strong conclusions, pretend there are just two periods—the present and the future—where the "future" is about one hundred years from now. For the purposes at hand, I am about to conduct a gigantic macroeconomic cost–benefit exercise trading off less present consumption from greenhouse gas abatement for more future consumption from mitigating the bad effects a century hence of global warming. Technically speaking, the possibility of extreme left-tail values of g occurring with small positive probability is outside this marginalist framework and requires us to go back to the fundamentals of expected utility theory that lie behind cost–benefit analysis under uncertainty, but I will cross that bridge when I come to it later and the take-away message will turn out to be similar anyway.

Of course such an incredible oversimplification of the economics of climate change ignores or distorts truly monumental chunks of reality. As just one example among many, a very important part of the global warming story concerns the huge stock–flow lags and enormous built-in inertias from having such a long pipeline between greenhouse gas emissions and ultimate temperature changes. This built-in inertia causes ΔT to continue to rise to levels above $E[\Delta T] \approx 3°\text{C}$ after a century from now (and also causes the tail probabilities of *very* high temperatures $\Delta T > 6°\text{C}$ to be relatively much bigger two centuries from now than one century from now) along any trajectory that does *not* stabilize atmospheric greenhouse gases at ≈ 550 ppm CO_2e—including the majority-opinion gradualist climate-policy-ramp trajectory. The ultimate high-temperature consequences of the huge inertial lag of ΔT to greenhouse gases already in the pipeline animate the *Stern Review* passion for severe curtailment of greenhouse gas emissions to begin soon, because at current flow rates we will attain a stock of 550 ppm of CO_2e within about a half-century and move (essentially irreversibly) thereafter beyond any hope to stabilize ultimate $E[\Delta T]$ at $\approx 3°\text{C}$. However, the point here is to put aside temporarily such details as the optimal consumption-smoothing profile of measures to slow greenhouse warming (and their inertial consequences) in favor of immediate transparency by focusing on the highly aggregated macroeconomic big picture of what is most essential in driving the *Stern Review* results, for which purpose focusing on a century hence is a good enough approximation.

Going right to the target here, my own rough point-guesstimate of what most economists might think are decent parameter values would be something like a "trio of twos": $\delta = 2$ percent, $g = 2$ percent (both on an annual basis), and $\eta = 2$ percent. For the sake of moving along, I am not going to try to defend the "trio of twos values" with a bunch of citations but instead I pretend for the time being that every critic of *Stern* thinks they are about right, so we can temporarily shelve this issue. Plugging these primitives into (1) makes "the" annual interest rate be $r = 6$ percent. Other reasonable—in my view—parameter combinations, say $\delta = 1$ percent, $g = 2$ percent, $\eta = 2.5$ (or even $\delta = 0$ percent, $g = 2$ percent, $\eta = 3$) also give $r = 6$ percent.

Concerning the rate of pure time preference, *Stern* follows a decidedly minority paternalistic view (which, however, includes a handful of distinguished economists) that for social discounting selects the lowest conceivable value $\delta \approx 0$ according to the a priori philosophical principle of treating all generations equally—irrespective of preferences for present over future utility that people seem to exhibit in their everyday savings and investment behavior. In a similar spirit of choosing extreme taste parameters, *Stern* selects as its base-case coefficient of relative risk aversion the value $\eta = 1$ that is the lowest lower bound of just about any economist's best-guess range. Some other taste-parameter values are considered in a halfhearted sensitivity analysis postscript to the original version of the *Review*, which is reported as if indicating robustness but I would interpret as more nearly the opposite because, no matter what spin is put on it, there is no escaping the impact of higher interest rates on undoing the *Review*'s extreme policy conclusions. With its preferred base-case parameter values $\delta = 0.1$ percent p.a., $g = 1.3$ percent p.a., $\eta = 1$, *Stern*'s discount rate from (1) is $r = 1.4$ percent. The present discounted value of a given global-warming loss from a century hence at the non-*Stern* annual interest rate of $r = 6$ percent is *one hundreth* of the present discounted value of the same loss at *Stern*'s annual interest rate of $r = 1.4$ percent. The disagreement over what interest rate to use for discounting is equivalent here in its impact to a disagreement about the estimated damage costs of global warming a hundred years hence of *two orders of magnitude*. Bingo!

If D is aggregate damages from global climate change and Y is GDP, then values of the ratio $\frac{D}{Y}$ a century from now (if nothing or very little is done to halt greenhouse gas emissions) are commonly taken to be somewhere in the range of about 0 percent to 3 percent. The *Stern Review* effectively uses $\frac{D}{Y} \approx 5$ percent as its base case. This high value is consistent with what an uncharitable critic might see as a philosophy of focusing on the gloomier outcomes in a heuristic-intuitive attempt to include extreme damages, because in *Stern*'s language "when we try to take due account of the upside risks and uncertainties, the probability-weighted costs look very large." Actually, the *Review* goes well beyond 5 percent in its multi-dimensional approach by making numerous literary and numerical allusions to the dark possibilities lurking in the tails of the distribution of possible outcomes (and then, as it were, rubbing salt in the wound of numerical calibration by noting how centrist it is actually being by *not* choosing much higher probability-weighted distant-future damages, which could be as big as $\frac{D}{Y} \approx 20$ percent–35 percent when one considers catastrophes that might materialize after 2105). *Stern* also estimates the annual costs of its ambitious abatement strategy as being equivalent to about 1 percent of GDP (which seems rather on the low side by maybe a factor of two or more, but that is not so relevant here).

The question for the *Stern Review* analysis then effectively becomes: is it worthwhile to sacrifice costs $C \approx 1$ percent of GDP now to remove damages $D \approx 5$ percent of GDP a century from now? With g and r being expressed on an annual basis, the benefit-over-cost ratio of such an investment would be $\frac{B}{C} = 5\exp(100(g - r))$. From (1), $r - g = (\eta - 1)g + \delta$, so that by picking the extreme values $\eta = 1$, $\delta = 0.1$ percent, *Stern* guarantees that the difference $r - g$ is *always* the miniscule amount $\delta = 0.1$ percent, no matter what value of g is chosen, which is really stacking the deck in favor of approving such kind of fractional GDP swaps across time. (The *Review* could have made life easier here by just rounding down a mere tenth of a percent by assuming $\delta = 0$, which along with $\eta = 1$ would make cost–benefit analysis *really* simple because a fixed fraction of GDP would then always be worth the same fixed fraction of GDP at *any* future time.) With *Stern*'s preferred parameter values, the benefit–cost ratio is $\frac{B}{C} = 4.5$ (close to the upper bound of $\frac{B}{C} = 5$ from assuming a zero rate of pure time preference)—a clear slam-dunk accept. The alternative non-*Stern* values $g = 2$ percent, $r = 6$ percent make $\frac{B}{C} = \frac{1}{10}$—a clear reject. This simple kind of exercise is what drives the *Stern Review* results and, in a nutshell, is what accounts for the difference with the more conventional analyses of its critics. The no-frills stripped-down variant of the Ramsey model I am using here is liable to a thousand and one legitimate questions and criticisms about its oversimplifications but, at the end of the day, I believe this exercise is highlighting fairly what really counts in the economics of climate change—the hidden discounting assumptions whose role tends to be more obscured than informed by the big IAMs.

For most economists, a major problem with *Stern*'s numbers is that people are not observed to behave as if they are operating with $\delta \approx 0$ and $\eta \approx 1$. To gauge the magnitude of the headache this presents for *Stern*'s taste-parameter values, consider the following thought experiment expressed in terms of the permanent income hypothesis in a deterministic setting. Suppose that on the margin an individual representing a long-lived dynasty faces a constant interest rate, *r*, and has a level of wealth, *W*, representing the capitalized value of future earnings plus initial holdings. Then permanent income is *rW* and an optimal consumption trajectory saves a constant amount *s* of permanent income. Plugging the implied balanced growth rate $g = sr$ into (1) and rearranging gives

$$s = \frac{r - \delta}{\eta r}. \tag{2}$$

With *Stern*'s preferred values $\delta \approx 0$, $\eta \approx 1$, equation (2) implies $s \approx 100$ percent irrespective of *r*—a *reductio ad absurdum*. In the economics of uncertainty, plausible values of the coefficient of relative risk aversion η are commonly taken to be somewhere between 1 and 4 (I use the geometric-average point estimate $\eta = 2$). A reader can plug favorite parameter values into (2) and back out implied values of δ. For me (and I suspect most economists), sensible savings rates in this and other variants of market-behavior-based thought experiments requires the rate of pure time preference to be significantly greater than zero (or at least if δ is chosen to be relatively small then η should be chosen to be relatively big). *Stern*'s worldview tends to blow off market-based observations and behavioral inferences as being (for a variety of reasons including market incompleteness) largely irrelevant to long-run discounting, which should instead be based primarily upon the "ethical" value $\delta \approx 0$ that *Stern* imposes on a priori grounds. Readers will have to make up their own minds about "ethical" values of preference parameters. While there may be *something* to *Stern*'s position about the limited relevance of market-based inferences for putting welfare weights on the utilities of one's great-grandchildren, and there might be *some* sporadic support for *Stern*'s preferred taste parameters scattered throughout the literature, I ultimately find such an extreme stance on the primacy of $\delta \approx 0$, $\eta \approx 1$ unconvincing when super-strong policy advice is so dependent upon nonconventional assumptions that go so strongly against mainstream economics.

3. *Puzzles and Ambiguities of Uncertain Discounting*

The most worrisome omission from any analysis based on the Ramsey approach (1) is uncertainty. As a first-pass informal cut at uncertainty, suppose we admit that we don't really know for sure whether *Stern* or *Stern*'s critics are right about the interest rate to use for discounting costs and benefits a hundred years or so from now. An important feature of interest rates under uncertainty is that they don't aggregate arithmetically into a simple certainty-equivalent interest rate. A $\frac{1}{2}$ chance of $r = 6$ percent and a $\frac{1}{2}$ chance of $r = 1.4$ percent are not at all the same thing as splitting the difference by selecting the average $r = 3.7$ percent. It is not discount *rates* that need to be averaged but discount *factors*. A $\frac{1}{2}$ chance of a discount factor of e^{-6} a century hence and a $\frac{1}{2}$ chance of a discount factor of $e^{-1.4}$ a century hence make an expected discount factor of $0.5e^{-6} + 0.5e^{-1.4}$ a century hence, which, when you do the math, is equivalent to an *effective* interest rate of $r = 2$ percent. According to this logic, the interest rate we should be using to discount a dollar of costs or benefits a century from now is in between the *Stern* value of $r = 1.4$ percent and the more conventional value of $r = 6$ percent, but with the above numbers it is a lot closer to the *Stern* value and is not anywhere near the arithmetic average of $r = 3.7$ percent. More generally here, if there is a subjective probability p_i that discount rate r_i is the correct rate to use, then the *effective* discount rate for time *t* is

$$r(t) = -\frac{\ln \sum p_i e^{-r_i t}}{t}, \tag{3}$$

which declines monotonically over time from the expected interest rate $r(0) = \sum p_i r_i$ to an asymptotic limit of $r(\infty) = \min_i \{r_i\}$. The moral of this story is that the *Stern* value may end up being more right than wrong when full accounting is made for the uncertainty of the discount rate itself, which arguably is the most important uncertainty of all in the economics of climate change. The very same force of compound interest that makes costs and benefits a century from now seem relatively insignificant, and that additionally creates the "majority tilt" of a pain-postponing climate policy ramp of emissions reductions starting from a low gradual base, also forces us to recognize the logic that over such long periods we should be using interest rates at the lower end of the spectrum of possible values.

In the certain world of the Ramsey deterministic formula (1), there is no distinction among rates of return on various assets and r is just *the* economywide rate of return on capital or, more succinctly, *the* interest rate. In nondeterministic reality, there are many rates of return out there and they differ considerably. The point has already been established that it makes a tremendous difference for long time periods of a century or more what interest rate is used for discounting. To understand better which discount rate to use, we need to enrich the Ramsey model by formally introducing uncertainty, which allows us at least to distinguish between rates of return on capital from two fundamentally different sorts of investments: a risky economywide rate of return applicable to investments that have payoff characteristics parallel to the economy itself and a risk free rate of return applicable to investments whose payoffs are orthogonal to the economy as a whole. After that, we need to decide which of these two rates is more appropriate for discounting costs and benefits of mitigating climate change. Then we need to plug in numbers and see what happens. The simplest formal way to begin this process is by making the growth rate be a random variable.

Continuing here in the spirit of being simple, suppose that the growth rate g in any given year is i.i.d. normal with *known* mean μ and *known* variance σ^2. (The fact that μ and σ^2 are known will later become significant when we inquire what happens under greenhouse warming when μ and σ^2 are modeled as *not* known.) With $g \sim N(\mu,\sigma^2)$, the Ramsey formula (1) becomes

$$r^f = \delta + \eta\mu - \frac{1}{2}\eta^2\sigma^2, \tag{4}$$

where r^f in equation (4) denotes the *risk free* interest rate. The introduction of uncertainty also allows consideration of a risky asset with a different rate of return. Following the asset-pricing expository literature, suppose we model comprehensive or representative equity at a high level of abstraction as being a claim on the consumption dividend produced by the macroeconomy itself. Suppose this abstract macroeconomy is represented by a dynamic stochastic general equilibrium in the Lucas–Mehra–Prescott fruit-tree model. Let the random variable R^e be the gross arithmetic return on equity while $r^e = \ln R^e$ is the more familiar geometric rate of return on equity. When g is i.i.d. $N(\mu,\sigma^2)$ in this fruit-tree economy, the equity risk premium over the safe rate then reduces to the well-known expression

$$\overline{r}^e - r^f = \eta\sigma^2, \tag{5}$$

where $\overline{r}^e$ is defined by the oblique-looking expected-value formula $\overline{r}^e \equiv \ln E[R^e]$, which is close enough to $E[r^e]$ to make them interchangeable for my purposes here.[4] Combining (5) with (4) gives the average return on equity as

$$\overline{r}^e = \delta + \eta\mu - \frac{1}{2}\eta^2\sigma^2 + \eta\sigma^2. \tag{6}$$

[4] The formulas (4) and (5) are explained in most graduate-level textbooks covering asset pricing.

Extending the previous "trio of twos" parameter values to a not-implausible knee-jerk "quartet of twos" $\delta = 2$ percent, $\eta = 2$, $E[g] = 2$ percent, $\sigma[g] = 2$ percent (on an annual basis, for long time series) makes very little difference on the risk free rate because now $r^f = 5.9$ percent in (4) instead of the previous value for "the" interest rate of $r = 6$ percent in (1). The corresponding equity premium from (5) is $\overline{r}^e - r^f = 0.1$ percent and the average return on equity from (6) is $\overline{r}^e = 6$ percent. The actual empirical numbers are closer to $r^f \approx 1$ percent, $\overline{r}^e - r^f \approx 6$ percent, $\overline{r}^e \approx 7$ percent. (The calibration $r^f \approx 1$ percent refers to short-term treasury bills, while $\overline{r}^e \approx 7$ percent refers to overall returns on comprehensive indexes of publicly traded shares of common stocks, but I don't think the numbers would be fundamentally different for other empirical measures of returns from investments for the economy as a whole.) So with the not-implausible "quartet of twos" parameter values the theory does a decent job of predicting the average return on equity but fails miserably on the risk free rate and the equity premium—thereby giving rise to the notorious "risk free rate puzzle" and the even more notorious "equity premium puzzle."

What does all of this have to do with the economics of climate change? Well, a lot actually. But before getting into the relationship between the asset-return puzzles and the economics of climate change, we need to put the puzzling numerical mismatches temporarily aside in favor of first asking a fundamental prenumerical question: *in principle* (leaving aside their correct numerical values) should we be using the risk free rate or the risky economywide rate of return for discounting costs and benefits of climate change?

The issue of which rate of return to choose (as between r^f and $\overline{r}^e$) for discounting a project comes down to the extent to which the payoffs from the project are proportional to or independent from returns to investments for the economy as a whole. In the oversimplified two-period formulation here, a project to mitigate the effects of global warming incurs consumption costs in the present period by curtailing CO_2e emissions, investing in costly new technologies, and so forth, but consumption in the future period is increased by having reduced the detrimental impacts at that time from greenhouse warming. The payoff is the extra consumption available in the distant-future period. Suppose that the correlation coefficient between the increased output of the project and returns to the economy as a whole is β. An investment beta is intended to represent a correlation coefficient that applies to discount *factors* as contrasted with discount *rates* (i.e., here β is the correlation between the investment payoff and R^e, not r^e). It then follows from essentially the same considerations as went into deriving formula (3) that the relevant interest rate for discounting costs and benefits at time t here is

$$(7)\quad r(t) = -\frac{\ln[\beta \exp(-\overline{r}^e t) + (1-\beta)\exp(-r^f t)]}{t},$$

which declines monotonically over time from $r(0) = \beta\overline{r}^e + (1-\beta)r^f$ to an asymptotic limit of $r(\infty) = r^f$. So the question here becomes: what is the right β for the kinds of projects that the *Stern Review* has in mind for mitigating global warming?

Overall damages from climate change are modeled in most IAMs, including the PAGE model that crunches some numbers for the *Review*, as a pure production externality equivalent to losing output via a particular subaggregator equation of the multiplicative form

$$(8)\qquad D(t) = Y^*(t) - Y(t) = f(\Delta T(t))Y^*(t),$$

where t is time, D is the total damages of greenhouse warming, ΔT is atmospheric temperature relative to the base period, Y^* is potential GDP (or NDP, no distinction being made here) in the absence of any greenhouse warming, and Y is actual GDP with greenhouse warming. The standard functional form actually chosen in most IAMs is $f(\Delta T) = k(\Delta T)^{\gamma}$ for some coefficients γ and k, where typically $\gamma = 2$ (quadratic loss in temperature change). The parameter k is usually calibrated so as to make $\frac{D}{Y}$ a century hence under mild or no abatement (with $\Delta T \approx 2.5$°C) be somewhere between approximately 0 percent and approximately 5 percent (depending on who is doing the calibrating). There is no question here about the value of beta implicit in the multiplicative formulation (8): it is one! Therefore, by the very logic of the IAM used by the *Stern Review* itself, the interest rate for discounting costs and benefits should be the returns to the economy as a whole, $\overline{r}^e$. This still leaves open the question of which rate to use for $\overline{r}^e$—the empirical returns on a broad index of publicly traded shares of stocks of about 7 percent (representing economy-wide average returns and used, e.g., by the Congressional Budget Office for evaluating U.S. government projects) or the value of 6 percent predicted by formula (5) from my non-*Stern* "quartet of twos" parameter values—but the discrepancy between 6 percent and 7 percent is insignificant for purposes here. Whatever number is used for $\overline{r}^e$, if it in any reasonable way represents the returns to the economy as a whole then it will completely undo the *Review* conclusions about drastic consumption smoothing and bring the results back to the much more moderate take-it-more-slowly climate-policy-ramp time profile of emissions reductions advocated by most mainstream critics of *Stern*.

This important dispute about what interest rate to use for discounting costs and benefits of mitigating greenhouse warming duplicates the same debate about the same subject more than a decade ago between William R. Cline and Nordhaus, two early pioneers of modeling the economic effects of climate change.[5]

[5] Cline (1992), Nordhaus (1994). See also the later studies and reflections on discounting for climate change contained in Paul R. Portney and John P. Weyant (1999).

Like *Stern*, the essentially identical earlier formulation of Cline used parameter values that made the Ramsey formula (5) deliver a low interest rate—in Cline's case the assumed parameter values were $\delta = 0$ percent, $\eta = 1.5$, $g = 1$ percent, which combined to make *the* interest rate be $r = 1.5$ percent per year. Also like *Stern*, the strong activist conclusions of Cline's analysis fifteen years earlier traced back to the very low discount rate being used. Furthermore, Cline and *Stern* are soulmates in their *cri de coeur* justifying $\delta \approx 0$ by relying mostly on a priori philosopher-king ethical judgements about the immorality of treating future generations differently from the current generation—instead of trying to back out what possibly more representative members of society than either Cline or *Stern* might be revealing from their behavior is *their* implicit rate of pure time preference. An enormously important part of the "discipline" of economics is supposed to be that economists understand the difference between their own personal preferences for apples over oranges and the preferences of others for apples over oranges. Inferring society's revealed preference value of δ is not an easy task in any event (here for purposes of long-term discounting, no less), but at least a good-faith effort at such an inference might have gone some way towards convincing the public that the economists doing the studies are not drawing conclusions primarily from imposing their own value judgements on the rest of the world.

In part because Cline's results, and where they were coming from, were more transparent (largely from not being buried within a big mysterious IAM, which was not yet readily available around 1990), his study attempted to seize the analytical high ground by emphasizing that an assumed annual interest rate of $r = 1.5$ percent is calibration-consistent with the real return on relatively safe U.S. Treasury bills historically being about 1 percent or so per annum. Missing from Cline's reasoning was a serious discussion of the implications of risk and of payoff correlations for the choice of a discount rate that might justify using $r = r^f$. Nordhaus, whose careful pragmatic modeling throughout his DICE series of IAMs has long set a standard in this arena, argued forcefully over a decade ago that the risk free interest rate should not be used for discounting costs and benefits of climate change. In this argument, Nordhaus was following Robert Lind who, in a comprehensive summary of an influential book he edited in 1984 entitled *Discounting for Time and Risk in Energy Policy,* concluded that "unless there is substantial evidence to the contrary, the returns associated with public projects should be assumed to be highly correlated with returns to the economy as a whole" (p. 77).

All of this having been said, there was never any deep economic rationale in the first place for damages from greenhouse gas warming being modeled as entering utility functions through the particular reduced-form route of being a pure production externality that substitutes perfectly with output according to the multiplicative subaggregator function (8). It was more due to an historical accident of stumbling upon a simple understandable analytical form whose parameters could be conveniently adjusted to match various scenarios than the result of serious thought about whether damages from global warming are better specified as multiplicative or additive with GDP, or even entering the utility function as a direct argument (rather than substituting one-for-one with economic output)—all of which would have been seen as a secondary issue. So, with the benefit of hindsight, let us now ask: Is there any economic rationale by which greenhouse-warming damages are as much uncorrelated as they are correlated with aggregate economic activity? The answer, when you think about it, is yes. No one has ever tried to argue that the effects of global warming will be evenly spread among regions of the world or sectors of the economy. The parts of an economy likely to be most impacted by global warming involve its "outdoor" aspects (broadly defined) like agriculture, coastal recreational areas, and natural landscapes (including the existence value of ecosystems, species, and so forth). Climate-affected "outdoor" activities may be differently impacted by greenhouse warming than "indoor" economic activities constituting the bulk of the economy, which are largely going to be dominated by the unknown future growth rate of labor-augmenting technological progress. Instances of changes in "outdoor" activities under global warming include what happens to tropical agriculture, losing significant parts of Bangladesh (or Florida) to rising sea levels, the "consumption" of an altered natural world that is a direct argument in the utility function, and so forth. These kinds of changes, which include the existence value of natural environments, are presumably not highly correlated with technological progress in computing power, furniture making, or better pharmaceuticals a century from now. The relevant share of the "outdoor" subset of the economy in investment-beta calculations might be disproportionately large because it is disproportionately largely impacted by greenhouse warming. Furthermore, it might plausibly be argued that the high income elasticity of environmental awareness will make for a high existence value of unaltered natural habitats when per capita incomes have increased ten-fold over the course of a century or more.

What happens to the discount rate for climate-change investments naturally depends on the actual value of β that is assumed. If $\beta = 0$ in (7), then $r(t) = r^f = 1$ percent. If $\beta = 1$ in (7), then $r(t) = \overline{r}^e = 7$ percent. The more interesting question concerns what happens to $r(t)$ for in-between values of β. Suppose for the sake of argument we split the difference and imagine that the disproportionate impact of climate change on generalized-land-usage "outdoor" activities of the economy warrants an overall correlation coefficient of, say, $\beta \approx 0.5$. With $\beta \approx 0.5$ in (7), the relevant interest rate for a century from now becomes $r(100) = 1.7$ pecent, which is close to Stern's $r = 1.4$ percent or Cline's earlier $r = 1.5$ percent. In this case investments for mitigating global climate change become attractive as an insurance policy that secures food supplies, preserves coastal areas, and maintains natural environments in a world where future aggregate growth rates are uncertain. I am not trying to defend this particular formulation or the particular value $\beta = 0.5$. Rather, the moral of this story is that the nature of the impacts of climate change determine whether we should end up closer to using the risk-free rate or the economywide return on capital—and there are plenty of stories suggesting that the relevant investment beta here is significantly less than one. When the overall discount factor is a combination of more primitive discount factors (as is the case here when the correlation coefficient β is some midrange value between zero and one), the risk free interest rate, which is close to the *Stern* interest rate, then may well end up being more right than wrong. Over a time horizon of a century or so, this "midrange β effect," which is not implausible when one considers the highly-uneven impacts of greenhouse warming on the different regions and sectors of the world economy, can be a strong factor in lowering discounting rates significantly—from the same underlying analytical source as the force of compound interest and the logic of the climate-policy ramp. Remarkably, the big IAMs with their casually built-in specification of $\beta \approx 1$ obscure rather than clarify the critical role in climate-change analysis of assumptions about investment betas.

Next, suppose we try to repeat the above numerical exercise but in place of the empirical values $r^f = 1$ percent, $\bar{r}^e = 7$ percent, we use the values predicted by the theoretical formulas via assuming the quartet of twos parameter values, which then implies $r^f = 5.9$ percent from (4) and $\bar{r}^e = 6$ percent from (6). Because the equity premium predicted from (5) is a miniscule 0.1 percent, there is essentially no difference in this case between $r^f = 5.9$ percent and $\bar{r}^e = 6$ percent. The relevant discounting rate $r(t)$ from (7) then lies between 5.9 percent and 6 percent independent of the assumed value of β. When 5.9 percent $\leq r(t) \leq 6$ percent, the *Review* conclusions are again undone and the more orthodox mainstream policies of moderate greenhouse gas slowing in the near future come back. The practical question of what interest rate to use for discounting costs and benefits of climate change thus becomes intertwined with the interpretation of the equity premium and risk free rate puzzles. It is a measure of how deep and serious these puzzles are that even after thousands of articles there is still no agreed-upon resolution of them. If we use numbers that resolve the puzzles in the descriptive direction, then r is sensitive to β and $r \approx 1.7$ percent for $\beta = 0.5$. If we use numbers that resolve the asset-return puzzles in the prescriptive direction, then $r \approx 6$ percent independent of β. And, to whip a horse long dead, it makes a huge difference to the economics of climate change whether $r \approx 1.7$ percent or $r \approx 6$ percent.

Critics of the *Stern Review* are fond of pointing out that $\delta \approx 0$, $\eta \approx 1$ is inconsistent with observed economic behavior, especially savings behavior. While this is true, and it is a genuine problem for *Stern*, it is just the tip of an iceberg that threatens all such formulations—not just *Stern*'s. The biggest and most troubling disconnect between the prescriptive numbers that theory says we should be using for discounting and the descriptive discount-rate numbers that are actually out there concerns the asset-return puzzles. These puzzles very strongly suggest that something fundamental is amiss in the paradigm framework for pricing assets and deriving the rates of return that we are relying upon to produce discount rates for evaluating new investment opportunities. For example, perhaps the taste parameters δ and η that we are commonly using (here $\delta = 2$ percent p.a. and $\eta = 2$) are wrong. If we treat (4) and (5) as two equations in two unknowns (δ and η), we can then invert the two equations to back out the hypothetical values $\hat{\delta}$ and $\hat{\eta}$ that would "explain" the stylized-fact empirical observation that $r^f \approx 1$ percent and $\overline{r}^e - r^f \approx 6$ percent. When this is done (for $\mu = 2$ percent, $\sigma = 2$ percent), it produces the mega-puzzle that the estimated rate of pure time preference is $\hat{\delta} \approx 151$ percent per year and the coefficient of relative risk aversion is $\hat{\eta} \approx 150$. One does not know whether to laugh or to cry at the prospect of what the *Stern Review* IAM might end up recommending as its preferred policy for climate change in its number-crunching simulations if the parameter values $\hat{\delta} \approx 151$ percent, $\hat{\eta} \approx 150$ were fed into PAGE. So much for the fantasy that values of the taste parameters δ and η should be chosen to be consistent with the revealed-preference observed stylized facts of economic behavior!

At the end of the day, where do these dizzying and disconcerting numerical exercises leave us with respect to the economics of climate change? One inescapably strong conclusion is that the emissions reductions that go along with optimal growth under endogenous climate change are extraordinarily sensitive to the interest rate that has implicitly been built into whatever model is being used for the analysis. The present discounted value of a future cost (or benefit) is the product of an imposed discount factor times the projected future cost (or benefit). Trying to forecast costs and benefits of climate change scenarios a hundred years or so from now is more the art of inspired guesstimating by analogy than a science (imagine forecasting today's world a century ago). But in my opinion the unsure prediction of future costs and benefits of climate change a century or two hence is overshadowed by the unsure interest rate to use in the discount factor, which makes the discount factor more uncertain than predicted costs (or benefits) of climate change by about an order of magnitude. Of the two multiplicands in the product of a discount factor times an expected cost (or benefit), empirically it is the discount-factor uncertainty that looms much larger in practice for analyzing climate-change-affected events a century or so from now.

4. *Uncertainty Tends to Matter Much More than Risk*

If the conclusion from the last section—that what to do about global warming depends overwhelmingly on the imposed interest rate—is seen as disappointing, then a second conclusion is likely to seem downright unnerving. As noted, the choice of appropriate discount rate is itself extraordinarily sensitive to seemingly arcane modeling details like the value of the climate-change investment beta and how the asset-return puzzles are resolved. One interpretation of the asset-return puzzles, which could also have some relevance for the economics of climate change, is the idea that investors are disproportionately afraid of rare disasters. These rare disasters are not fully reflected in the available data samples that, being limited, are naturally deficient in coverage. Besides, even if we had an infinite time series of past observations, they are of restricted relevance in an evolving world whose features are always changing and whose past never fully repeats itself. With this interpretation of the puzzles, people are willing to pay high premiums for relatively safe stores of value that might represent "catastrophe insurance" against out-of-sample or newly evolved rare disasters.[6] Such an ongoing catastrophe-insurance effect could readily explain why observed r^f is so low relative to the observed past average of realized r^e.

There is little doubt that the worst-case scenarios of global-warming catastrophes are genuinely frightening. The *Stern Review* goes over several of these highly unlikely, poorly understood threshold-crossing disasters associated with abrupt large-scale irreversible changes in the climate system:

[6] The theme of catastrophe insurance and the underlying motivation for the treatment of structural uncertainty as tail thickening of posterior-predictive distributions is developed in Martin L. Weitzman (forthcoming).

sudden collapse of the Greenland and West Antarctica ice sheets, weakening or even reversal of thermohaline circulations that might radically affect such things as the Gulf Stream and European climate, runaway climate-sensitivity amplification of global warming due to positive-reinforcing multiplier feedbacks (including, but not limited to, loss of polar albedo, weakened carbon sinks, and rapid releases of methane from the thawing of arctic permafrost). More gradual but still very serious examples of uncertain climate-change effects are: sea-level dynamics, drowned coastlines of unknown magnitude, very different and possibly extreme weather patterns including droughts and floods, ecosystem destruction, mass species extinctions, big changes in worldwide precipitation patterns and distribution of fresh water, tropical-crop failures, large-scale migrations of human populations, humidity-nourished contagious diseases—and the list goes on and on.

Translated into the language of the simple model used here, such rare disasters are far out in the right tail of very high ΔT, which corresponds to being far out in the left tail of the consumption-growth random variable g. The probability distribution of long-run ΔT is disturbingly spread apart, largely because of structural-parameter uncertainty about the unknown "climate sensitivity" multiplier that amplifies greenhouse gas concentrations into ultimate steady-state greenhouse warming. The recently released *Fourth Assessment Report of the IPCC* (2007) predicts for one hundred years from now a mean temperature change of further planetary warming (from averaging six "equally sound" marker scenarios) of $E[\Delta T] = 2.8^\circ C$ with a thick-tailed upper-end standard deviation $\approx 1.6^\circ C$ (table SPM-3). This means the probability that $\Delta T > 4.5^\circ C$ is approximately 15 percent and the probability of $\Delta T > 6^\circ C$ is very roughly about 3 percent. IPCC does not extend its projections beyond 2105 on the basis that predictions into the twenty-second century are too uncertain, but it seems unavoidable that the reduced-form probability of $\Delta T > 6^\circ C$ increases substantially above 3 percent *after* the next century just from the enormous inertial lags for what by then will be in the climate-change pipeline. Societies and ecosystems whose average temperature has changed in the course of a century or so by $\Delta T > 6^\circ C$ (for U.S. readers: $\Delta 6^\circ C \approx \Delta 11^\circ F$) are located in the *terra incognita* of what any honest economic modeler would have to admit is a planet Earth reconfigured as science fiction, since such high temperatures have not existed for some tens of millions of years.

The idea behind analyzing climate-change projects by converting future costs and benefits into present discounted values is that society has alternative investment opportunities, whose proxy rate of return is the discount rate, representing alternative capital-accumulation opportunities throughout the rest of the economy that would compensate us for the economic losses suffered from climate change. Human-capital investments in education or public health have consistently been found to have high rates of return, arguably far greater than 10 percent for less-developed countries and regions. More mundane examples of alternatives to CO_2e mitigation from middle-of-the-probability-distribution mild warming might include accumulating air conditioners to counter high temperatures or erecting sea walls to keep the rising ocean out of coastal cities. Such alterative investments compensate mostly for potential loss of "indoor" consumption and they tend to be a lot less expensive than wholesale abatement of greenhouse gases. The real problem is in the tails and it mostly concerns "outdoor" consumption. If the definition of consumption is broadened (as it should be) to include nonmarket enjoyment of the natural environment—like habitats, ecosystems, and species—then it is difficult to imagine what the compensating investments are for which we should now be saving more as an alternative that might substitute for holding down ΔT directly. With roughly 3 percent IPCC-4 probability, we will "consume" a *terra incognita* biosphere within a hundred years whose mass species extinctions, radical alterations of natural environments, and other extreme outdoor consequences of a different planet will have been triggered by a geologically-instantaneous temperature change that is significantly larger than what separates us now from past ice ages.

In the rest of this paper, marginal analysis is set aside and g stands for the unknown growth rate of a comprehensive future "consumption" that includes the consumption of natural environments, ecosystems, species, and the like. The cost of low-g disasters from high-ΔT scenarios more properly constitutes uncertainty in the sense of Knight or Keynes than risk, because the scale and probability of these disasters are both unknown. Not only is it very difficult to estimate tail probabilities of high-ΔT outcomes—due, ultimately, to the underlying sampling-theory principle that the rarer is an event the more unsure is our estimate of its probability—but translating this into g-equivalent economic-damage units introduces enormous further fuzziness, especially when g includes existence values of natural habitats. With an evolutionary stochastic process like global climate change, the world is not standing still long enough for us to accumulate the relevant information to accurately assess tail probabilities. The net result is thicker left tails for the distribution of g under dynamically evolving global climate change than we are accustomed to dealing with in our much more familiar dynamic stochastic general equilibrium macro models, which in practice are based upon the stationary thin-tailed stochastic processes that we use to model a rational expectations equilibrium whose structure is (supposedly) fully known and understood.

Every cost–benefit analysis is an exercise in subjective uncertainty. If, as the *Stern Review* puts it, "climate change is the greatest externality the world has ever seen," then a cost–benefit calculation of what to do about it is the greatest exercise in Bayesian decision theory that we economists have ever performed. Formally, of course, cost–benefit analysis can deal with uncertainty—by taking expected values, relying on expected-utility theory, accounting for risk aversion, and using all of the other, by now familiar, paraphernalia of the modern theory of the economics of uncertainty. In principle, it does not matter whether the probabilities that show up in our cost–benefit calculations are objective or subjective because the mathematical formulas are the same for either case. But in lumping together objective and subjective uncertainties and thereby obscuring their distinction—to the extent that a graduate student today hardly knows, or even cares, what kinds of probabilities are legitimate to plug into a rational expectations equilibrium and what kinds of probabilities are illegitimate for such purposes—I think that contemporary macroeconomics goes too far and leads to a mindset that too easily identifies probability (and "economic science") with exercises in calibration to sample frequencies from past data.

I do not propose to rehash here the ages-old, never-resolved foundational controversy about whether probabilities are better conceptualized on the most fundamental level as objective frequencies or subjective beliefs. Personally, I do not think there exists a pure case of either extreme pole, but rather there is a continuum of situations with some being closer for practical modeling purposes to the objective pole and others being closer for practical modeling purposes to the subjective pole. Here I just want to point out that if something like radioactive decay is close to being a pure case of objective frequencies, then climate change, and especially the *economics* of climate change, is as close to being a pure case of modeling probabilities by subjective judgements as we economists are ever likely to encounter in practice. To paraphrase the language of the *Stern Review* yet again, the economics of climate change is the greatest application of subjective uncertainty theory the world has ever seen.

To the extent that it makes any sense at all to think in terms of some approximately bell-shaped meta-distribution of growth rates g that is out there, the part of the probability distribution that corresponds most closely to objective-frequency risk is in its body around the middle because, from previous experience, past observations, plausible extrapolations, and maybe even the law of large numbers, we have at least some modicum of confidence in being able to construct a reasonable approximation of the central regions of the probability distribution. As we move toward probabilities in the tails of the g distribution, however, we are increasingly moving into the unknown territory of subjective uncertainty where our probability estimates of the probability distributions themselves becomes increasingly diffuse because the frequencies of rare events in the tails cannot be pinned down by previous experiences, past observations, or computer simulations. The upshot of this uncertainty about uncertainties is that the reduced-form probability distribution of g (after integrating out the probabilities of probabilities)—which is a reduced form for the economics of climate change in the sense that g here is the growth rate of comprehensive consumption that includes the natural environment—has a thick left tail. The exact thickness of this left tail of g depends not only upon how bad an environmental catastrophe global warming might induce and with what probabilities, but also upon how imprecise are our probability estimates of the probabilities of those bad catastrophes. Uneasiness about projecting uncertain uncertainties prevents IPCC and most economic analyses from taking a stand on the increasing—and increasingly diffuse—probabilities of extreme temperatures after the year 2105, which hardly eliminates the underlying problem. Mitigating the future consequences of greenhouse warming does not just shift the center of the distribution of g to the right but, far more importantly in this context, it thins the left tail of the distribution as well.

The thickened tails of the reduced form of the distribution of g that are an inevitable consequence of taking expectations of expectations can have surprisingly strong effects on cost–benefit calculations by lowering significantly expected utility and raising significantly expected marginal utility. To get a sense of just how strong the effect can be of tails thickened by having structural parameters that we do not know but whose values must be inferred indirectly from limited experience—and therefore a sense of how much we could be missing in our economic analysis by ignoring the *terra incognita* of the greenhouse-warming extremes—consider this prosaic example. Suppose that in the good old days before we understood human-induced climate change we were sure that $g \sim N(\mu,\sigma^2)$, where we somehow knew that $\mu = 2$ percent and $\sigma = 2$ percent. Normalize current marginal utility to be unity. Then from using the familiar formula for the expectation of a lognormally distributed random variable, the expected marginal utility of an extra sure unit of consumption in the pre-climate-change era would have been $EMU = E[\exp(-\eta g)] = \exp(-\eta\mu + \frac{1}{2}\eta^2\sigma^2)$. (It is precisely this kind of calculation that lies behind the risk free rate and equity-premium formulas (4) and (5).)

Imagine next that the possibility of greenhouse warming has now made us unsure about μ and σ. Let us preliminarily model this greenhouse-warming-induced uncertain situation, where we don't know the true values of μ and σ because of limited experience with climate change, *as if* we are limited because we only have data from some finite number n of past observations (or finite simulation outcomes from the data generating process of some model) and we run a regression to estimate μ and σ. For simplicity, suppose further that the point estimates $\hat{\mu} = 2$ percent and $\hat{\sigma} = 2$ percent from this regression just so happen to be the very same numbers as the presumed-known population parameters for the normal distribution before climate change was understood to be a possibility. Then the reduced-form situation is *as if* g is distributed as a Student-*t* distribution with $n - 1$ degrees of freedom. The Student-*t* here has the same mean as the normal and for large *n* has almost the same standard deviation, but if you look closely with a magnifying glass its tails are naturally thickened due to the "true" values of the structural parameters μ and σ being uncertain. This kind of structural uncertainty about the parameters of the probability distribution spreads apart the reduced-form ("predictive posterior" in Bayesian jargon) distribution of *g*, an effect that is especially pronounced in the thickened tails because they are especially difficult to learn about. If we now calculate the expected marginal utility of an extra sure unit of consumption using this Student-*t* distribution (which is a natural manifestation of limited experience or limited information), then $EMU = E[\exp(-\eta g)] = +\infty$, which is mathematically equivalent to the fact that the moment generating function of a Student-*t* distribution is infinite. The bombshell fact that $EMU = +\infty$ (as soon as we admit that we don't know the underlying stochastic structure, and therefore parameters must be estimated) changes the rules of the game. Such a mechanism, for example, explains the asset-return puzzles for reasonable values of δ and η as being due to a fear of relatively rare tail disasters that is theoretically difficult or impossible to eliminate when the underlying tail-structure remains uncertain. The fact that under structural uncertainty $EMU = +\infty$ represents a mathematically generic result not limited to isoelastic utility or the normal parent distribution and Student-*t* child distribution of the example. I claim this general result has significant economic repercussions which are not easily brushed aside, not least of all for cost–benefit analysis of climate change because such an effect in principle overshadows the discounting of far-future events.

There is a general point here and a particular application to the economics of climate change. The general point is that from experience alone one cannot acquire sufficiently accurate information about the probabilities of tail disasters to prevent the expected marginal utility of an extra sure unit of consumption from becoming unbounded for *any* utility function having everywhere-positive relative risk aversion, thereby effortlessly driving cost-benefit applications of expected utility theory. The degree to which this kind of "generalized precautionary principle" is relevant in a particular application must be decided on a case-by-case basis that depends upon the extent to which a priori knowledge in a particular case limits the extent of posterior-predictive tail thickening. In the particular application to the economics of climate change, where there is so obviously limited data and limited information about the global catastrophic reach of climate extremes for the case $\Delta T > 6°C$, to ignore or suppress the significance of rare tail disasters is to ignore or suppress what economic theory is telling us loudly and clearly is potentially the most important part of the analysis. While it is always fair game to challenge the assumptions of a model, when economic theory provides a generic result (like "free trade is Pareto optimal") the burden of proof is commonly taken as resting on whomever wants to overturn the theorem in a particular application. The take-away message here is that the burden of proof in the economics of climate change is presumptively upon whomever wants to model optimal-expected-utility growth under endogenous greenhouse warming *without* having structural uncertainty tending to matter much more than risk. Such a middle-of-the-distribution modeler needs to explain why the inescapably thickened tails of the posterior-predictive distribution, for which the thick left tail of *g* represents rare disasters under uncertain structure, is *not* the primary focus of attention and does *not* play the decisive role in the analysis.

5. *Climate Uncertainty and the Value of Information*

Because the *Stern Review* is imbued with the laudable moral imperative of not exposing future generations to the tribulations of global warming, it does not shy away from emphasizing (at least discursively in its "multidimensional" text) the possibilities of rare high-ΔT, low-g catastrophes from climate change. Indeed, reading between the lines of the report, one has the feeling that the immorality of relegating future generations to live under the shadow of the open-ended possibilities of uncertain large-scale changes in the climate system, when for a mere annuity cost of a percent or two (or at most three) of GDP each year we might have purchased an insurance policy on their behalf that avoided this scary uncertainty (or at least greatly reduced it), is a major underlying leitmotif of the *Review*. This feeling of guilt has no place to go analytically (under the conventional analytical confines adopted by the IAMs, including PAGE), so to speak, except to be subliminally channelled into choosing such low values of $\delta \approx 0$ and of $\eta \approx 1$ (and, secondarily, such high values of $\frac{D}{Y} \approx 5$ percent and low values of $\frac{C}{Y} \approx 1$ percent) as will operate through the back door of conventional economic analysis to weight present-discounted future damages high enough relative to present mitigation costs to make the IAM want to reduce substantially the disastrous possibilities. The *Review* puts it directly: "Averaging across possibilities conceals risks. The risks of outcomes much worse than expected are very real and they could be catastrophic. Policy on climate change is in large measure about reducing these risks. They cannot be fully eliminated, but they can be substantially reduced. Such a modeling framework has to take account of ethical judgements on the distribution of income and how to treat future generations."

The "ethical judgements" in the above quote about "how to treat future generations" is *Stern*-speak for picking $\delta \approx 0$, while the "ethical judgements on the distribution of income" is *Stern*-speak for picking $\eta \approx 1$. Such "ethical judgements" could appear to an uncharitable critic as if designed to justify the activist conclusions that are considered necessary to avoid the climate-change horror scenarios. In the context of its self-imposed "multidimensional" methodology, by choosing $\eta \approx 1$ the *Review* appears to be playing both sides of the street against the middle. On the one hand, it wants η to be as high as possible to reflect its tremendous humanitarian concern with distributional inequities *across space*, which would allow it to argue (informally, if passionately, in scattered prose and numerical examples) that the disproportionate negative impact of climate change on the world's poor (whose marginal utility is high because η in this story is implicitly large) calls for urgent action now to avoid future massive spatial redistribution of relative income from the poor to the rich. Simultaneously, the *Stern Review* wants to further exacerbate distributional inequities by redistributing income *across time* from the relatively poor present to the much-richer future a century or two from now (when standards of living are likely to be ten times higher) via—an uncharitable critic might suspect—choosing the lowest imaginable value $\eta \approx 1$ that might be used (along with $\delta \approx 0$) to reverse-engineer the really low r needed to prop up its technical case for immediate urgent action. The same contradiction about values of η shows itself in *Stern*'s heuristic justification for big values of probability-weighted $\frac{D}{Y}$ as being due, in effect, to a risk-averse high-curvature high-η utility function interacting with highly uncertain damages.

I think that rather than trying to go though the back door with unreasonably low values of δ and η (or, secondarily, "averaging across possibilities" by heuristically making business-as-usual $\frac{D}{Y} \geq 5$ percent instead of some smaller more plausible point estimate), it is much better to go directly through the front door with the legitimate concern that there is a chance, whose subjective probability is small but diffuse (thereby resulting in a dangerously thickened left tail of comprehensive consumption growth rates), that global warming may eventually cause disastrous temperatures and environmental catastrophes. If one accepts that global climate change is as likely an arena as any for a valid application of the general principle that thickened tails from uncertain structural parameters must dominate expected discounted utility calculations, then many hard questions need to be asked. What are early warning signs of impending runaway environmental disasters like melting ice sheets, thermohaline inversions, or just plain knowing beforehand that we are on a trajectory toward $\Delta T > 6°C$? How much would it cost to put in place the very best system of sensors that money can buy for detecting early-warning signals of impending climate catastrophes? How early might the warning from monitoring systems be before the full effects are felt? What could we do as an emergency response if we received such early-warning signals? Would last-ditch emergency geoengineering measures to ward off disaster by reversing the worst consequences of global warming be available in time to help? (Such emergency measures are likely to be so extreme as to be defensible only for an even more extreme environmental catastrophe in the making—perhaps they might include painting all human-made structures on the planet reflective white and creating a "Pinatubo effect" by seeding the upper atmosphere with metallic dust or aerosols.[7]) Could such last-ditch measures be made reversible by building in decay mechanisms, as with sulfate aerosols, while we then used the new information to *really* undertake draconian measures to cut greenhouse gas emissions drastically? Are there other aerosol precursors than sulfates with possibly better environmental properties? Can the public embrace the currently politically incorrect idea, which is a third rail few policymakers dare to touch and *Stern* doesn't even mention, that in the extremely unlikely event of a truly extraordinary unfolding disaster it might be a good emergency-backup plan to purposely geoengineer spaceship Earth to reverse previous inadvertent geoengineering from burning too much fossil fuel? Or, perhaps more to the point here, can anyone imagine how the public would *not* embrace such an idea and demand that we should do all in our power to avert a climate-change-induced catastrophe at some hypothetical future time when environmental disaster seems imminent? And how does the tail thickness of climate-change disasters compare with the tail thickness of aerosol geoengineering, or the tail thickness inherent in widespread nuclear power possibly going awry, or the tail thickness of massive sudden-release mishaps from buried-CO_2 sequestration because some remote "hypothetical" materialized?

I trust that a few readers may be able to think of more such questions about the real option value of waiting to gather information (and the empirical issue of what to do about it), some of which might hopefully be more grounded in reality than my highly speculative examples. Whether or not my particular hypothetical stories are realistic, these *kinds* of questions become relevant once the focus of the economics of climate change shifts from the middle range of the distribution of what might happen with ΔT at a IPCC-4 mean of $\approx 2.8°C$ a hundred years hence to thinking more about what might be in the tails with $\Delta T > 6°C$, which is just two IPCC-4 standard deviations out for a century from now, meaning a probability ≈ 3 percent (and presumably a yet higher probability *after* 2105). This thick tail is where most of the cost–benefit action may well be even if—or perhaps precisely because—our estimates of the probabilities involved are themselves so highly uncertain. Anything is *possible* in the tails of a nondogmatic distribution.

[7] On the feasible use of sulfate aerosol precursors to reverse global warming, see T. M. L. Wigley (2006).

("Nondogmatic" in Bayesian parlance just means that no event is ruled out a priori by having been assigned zero probability in the prior distribution.) A responsible policy approach neither dismisses the horror stories just because they are two standard deviations away from what is likely nor gets stampeded into overemphasizing false dichotomies as if we must make costly all-or-nothing investment decisions right now to avoid theoretically possible horrible outcomes in the distant future.

In my opinion, public policy on greenhouse warming needs desperately to steer a middle course, which is not yet there, for dealing with possible climate-change disasters. This middle course combines the gradualist climate-policy ramp of ever-tighter greenhouse gas reductions that comes from mainstream mid-probability-distribution analysis (under reasonable parameter values) with the option value of waiting for better information about the thick-tailed disasters. It takes seriously whether or not possibilities exist for finding out beforehand that we are on a runaway-climate trajectory and—without "leaving it all up to geoengineering"—confronts honestly the possible options of undertaking currently politically incorrect emergency measures if a worst-case nightmare trajectory happens to materialize. The overarching concern of such a middle course is to be constructive by having some semblance of a game plan for dealing realistically with what might conceivably be coming down the road. The point is to supplement mainstream economic analysis of climate change (and mainstream ramped-up mitigation policies for dealing with it) by putting serious research dollars into early detection of rare disasters and by beginning a major public dialogue about contingency planning for worst-case scenarios perhaps akin to the way Americans (at their best) might debate the pros and cons of an anti-ICBM early warning system. It may well turn out that the option value of waiting for better information about catastrophic tail events is negligible because early detection is impossible, or it is too expensive, or it comes too late (this is *Stern*'s line, and it might, or might not, happen to be true), or because nothing practical can be done about reversing greenhouse warming anyway—so we should stop stalling and start making serious down payments on catastrophe insurance by cutting CO_2e emissions drastically. But these are conclusions we need to reach empirically, rather than prejudging them initially. Instead of declaring immediate all-out war on greenhouse gas emissions as advocated by *Stern*, maybe we would do better by steadily but surely ramping up greenhouse gas cuts over the next decade or two while simultaneously investigating *seriously* the nature of the runaway-climate disasters in the thick tails and what might be done realistically about them should they start to materialize. We can always come back in ten or twenty years time and declare all-out war on global-warming emissions then—*if we then* think it is the best option among a better-studied reasonably considered portfolio of possible options.

Until we start seriously posing and trying to answer tough questions about rare global-warming catastrophes, we will not make real progress in dealing constructively with the nightmare scenarios and we will continue to cope with them inadequately by trying to shoehorn disaster policy into an either-or response category where it won't fit. The *Stern Review* has its heart in the right place—it is not nice for us to play the role of nature's grim reaper by bequeathing the enormously unsettling uncertainty of a very small, but essentially unknown (and perhaps unknowable), probability of a planet Earth that in hindsight we allowed to get wrecked on our watch. However, *Stern* does not follow through formally on this really unsettling part of the global warming equation (which a generous interpretation of its not-convincing economic analysis might say is the underlying motivation for its overall alarmist tone) except indirectly, by choosing $\delta \approx 0$, $\eta \approx 1$, $\frac{D}{Y} \approx 5$ percent, $\frac{C}{Y} \approx 1$ percent (which an ungenerous interpretation might say is reverse-engineering the drastic slowing measures that the *Review* wants to impose on greenhouse gas emissions to neutralize the nightmare scenarios). I don't mean to imply that there is some off-the-shelf turnkey consensus model of the economics of uncertain catastrophes that the *Stern Review* was negligent in not using or that such a model would (or should) provide ammunition for an excuse not to undertake serious action soon to slow greenhouse emissions. We just don't yet know and we need badly to find out. The overarching problem is that we lack a commonly accepted usable economic framework for dealing with these kinds of thick-tailed extreme disasters, whose probability distributions are inherently difficult to estimate (which is why the tails must be thick in the first place). But I think progress begins by recognizing that the hidden core meaning of *Stern vs. Critics* may be about *tails vs. middle* and about *catastrophe insurance vs. consumption smoothing*.

6. *Getting it Right for the Wrong Reasons?*

The *Stern Review* is a political document—in Keynes's phrase an essay in persuasion—at least as much as it is an economic analysis and, in fairness, it needs ultimately to be judged by both standards. To its great credit, the *Review* supports very strongly the politically unpalatable idea, which no democratic politician planning to remain in office anywhere wants to hear, that (however it is packaged and whatever spin is put on it) substantial carbon taxes must be levied because energy users need desperately to start confronting the expensive reality that burning carbon has a significant externality cost that ought to be taken into account by being charged full freight for doing it. (This is the most central "inconvenient truth" of all, which was conveniently ignored in Al Gore's award-winning film.) An entire chapter 22 in *Stern*, entitled "Creating a Global Price for Carbon," is devoted to this theme. As the *Review* puts it, "the pricing of carbon, implemented through tax, trading, or regulation," is "required for an effective global response" (p. xvii). One can only wish that U.S. political leaders might have the wisdom to understand and the courage to act upon the breathtakingly simple vision that steady pressure from the predictable presence of a high carbon price reflecting social costs (whether imposed directly through taxes or indirectly via tradable permits) would do more to unleash the decentralized power of capitalistic American inventive genius on the problem of researching, developing, and finally investing in economically efficient carbon-avoiding alternative technologies than all of the piecemeal command-and-control standards and patchwork subsidies making the rounds in Washington these days.

As we have seen, on the economic-analysis side the *Stern Review* predetermines the outcome in favor of strong immediate action to curtail greenhouse gas emissions by creating a very low value of $r \approx 1.4$ percent via the indirect route of picking point-estimate parameter values $\delta \approx 0$ and $\eta \approx 1$ that are more like theoretically reasoned extreme lower bounds than empirically plausible estimates of representative tastes. But we have also seen that a fair recognition of the truth that we are genuinely uncertain about what interest rate should be used to discount costs and benefits of climate changes a century from now brings discounting rates down from conventional values $r \approx 6-7$ percent to much lower values of perhaps $r \approx 2-4$ percent, which would create a more intermediate sense of urgency somewhere between what the *Stern Review* is advocating and the more modest measures to slow global warming advocated by many of its critics. The important remaining caveat is that such an intermediate position is still grounded in a conventional consumption-smoothing approach to the economic analysis of climate change that avoids formally confronting the issue of what to do about catastrophe insurance against the possibility of thick-tailed rare disasters, which from first principles of economic–statistical reasoning presumptively drive expected discounted utility outcomes.

In conclusion, I think the *Stern Review* deserves credit for effectively raising the level of public discourse—by increasing general awareness that climate change is a serious issue which should be taken seriously, by arguing cogently for what is effectively a global carbon tax as an essential component of any reasonable solution, by openly discussing adaptation to climate change (as well as mitigation), and by popularizing for a wider audience than economists the idea that economic analysis of costs and benefits might be a useful legitimate method for evaluating policies to mitigate global warming. I think also that *Stern* deserves some measure of credit for elevating to prominence the problem of genuine uncertainty concerning rare but dangerous climate disasters, about which decisions must be made but whose scale and probability cannot be known precisely. However, in my opinion, *Stern* deserves a measure of *dis*credit for giving readers an authoritative-looking impression that seemingly objective best-available-practice professional economic analysis robustly supports its conclusions, instead of more openly disclosing the full extent to which the *Review*'s radical policy recommendations depend upon controversial extreme assumptions and unconventional discount rates that most mainstream economists would consider much too low. I can't help but feel after reading the *Review* that its urgent tone of morality and alarm comes not from the chapter 6 aggregative economic modeling of climate-change impacts, which deservedly has drawn strong criticism from economists, but more from a not formally articulated fear of what might potentially be out there with greenhouse warming in (using my own ponderous terminology here to make sure my expression is exact) the inherently thickened left tail of the reduced-form posterior-predictive probability distribution of the growth rate of a comprehensive measure of consumption that includes the natural environment. I have argued that this inherently thickened left tail of g is an important aspect of the economics of climate change that every analyst—*Stern* and the critics of *Stern*—might do well to try to address more directly. History will judge whether the economic analysis of the *Stern Review* ended up being more wrong or more right and, if it was more right, whether as pure economic reasoning it was right for the right reasons or it was right for the wrong reasons.

References

Cline, William R. 1992. *The Economics of Global Warming*. Washington, D.C.: Institute for International Economics.

Dasgupta, Partha. 2007. "Commentary: The *Stern Review*'s Economics of Climate Change." *National Institute Economic Review*, 199: 4–7.

Hope, Chris. 2006. "The Marginal Impact of CO_2 from PAGE2002: An Integrated Assessment Model Incorporating the IPCC's Five Reasons for Concern." *Integrated Assessment*, 6(1): 19–56.

Intergovernmental Panel on Climate Change. 2007. *Climate Change 2007: The Physical Science Basis: Summary for Policy Makers*. http://www.ipcc.ch.

Kelly, David L., and Charles D. Kolstad. 1999. "Integrated Assessment Models for Climate Change Control." In *The International Yearbook of Environmental and Resource Economics: 1999/2000: A Survey of Current Issues*, ed. H. Folmer and T. Tietenberg. Cheltenham, U.K. and Northampton, Mass.: Elgar, 171–97.

Lind, Robert, ed. 1982. *Discounting for Time and Risk in Energy Policy.* Washington, D.C.: Resources for the Future.

Maddison, David. 2006. "Further Comments on the *Stern Review.*" http://www.economics.bham.ac.uk/maddison/Stern%20Comments.pdf.

Mendelsohn, Robert O. 2007. "A Critique of the Stern Report." *Regulation*, 29(4): 40–47.

Nordhaus, William D. 1994. *Managing the Global Commons: The Economics of Climate Change*. Cambridge, Mass. and London: MIT Press.

Nordhaus, William D. 2007. "A Review of the *Stern Review on the Economics of Climate Change*." *Journal of Economic Literature*, 45(3): 686–702.

Portney, Paul R., and John P. Weyant, eds. 1999. *Discounting and Intergenerational Equity*. Washington, D.C.: Resources for the Future.

Tol, Richard S. J., and Gary W. Yohe. 2006. "A Review of the *Stern Review*." *World Economics*, 7(4): 233–50.

Weitzman, Martin L. Forthcoming. "Subjective Expectations and Asset-Return Puzzles." *American Economic Review*.

Wigley, T. M. L. 2006. "A Combined Migration/Geoengineering Approach to Climate Stabilization." *Science*, 314(5798): 452–54.