The Stern Review of the Economics of Climate Change

Martin L. Weitzman* (13.02.07)

Book report commissioned by *JEL*. Preliminary. Liable to further changes. Reactions appreciated and comments welcomed.

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 anew about how to model (or at least how to envision) such fundamental concepts as risk, uncertainty, prediction, and discounting. There is nothing like being asked for a specific policy recommendation on a vivid actuality to put a little life into otherwise arcane matters of economic analysis. Beyond the issue of whether it is right or wrong in its conclusions, *The Stern Review of 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 both economic and political reasons 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 600+ pages, the *Stern Review* is comprehensive in its scope and ambitious in its aims. An economist wanting to get a feel for the basic issues of global climate change could do much worse than browse through this report, as it summarizes well the essentials. The *Review* contains much of value aside from 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 on line. Instead, at the risk of oversimplification I concentrate here on trying to distill

*Department of Economics, Harvard University, Cambridge, MA 02138 (e-mail: mweitzman@harvard.edu). For helpful detailed comments on earlier drafts of this paper, but without implicating them for its remaining defects, I am grateful to Olivier Blanchard, Richard Cooper, Stephen DeCanio, Roger Gordon, Lawrence Summers, and Hal Varian.

the *Review* down to what I think is its analytical essence as a piece of applied cost-benefit analysis, because it can be a bit difficult to see 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 greenhouse gas emissions because (and this quote captures well the tone of urgency about moving quickly to avoid catastrophic possibilities, which 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 twentieth century." Such a strong call to immediate decisive action is at odds with what most other economic analyses of climate change have concluded. By contrast, 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 gradually tighten considerably over time. The *Review* analysis, on the other hand, finds that "the benefits of strong, early action on climate change outweighs the costs" and calls for stabilizing greenhouse-gas atmospheric concentrations at $\approx$550 parts per million (ppm) of CO_2-equivalent (CO_2e). (The current level is $\approx$430 ppm CO_2e, compared with $\approx$280 ppm CO_2e before the Industrial Revolution.) This would make temperatures a hundred years from now be at $\Delta T \approx$1.5-2.5°C compared with pre-industrial-revolution values and would (hopefully) stabilize future temperatures permanently thereafter at $\Delta T \approx$3°C Instead, on the more-gradual majority optimal trajectory, $\Delta T \approx$2-3°C a century from now and temperatures continue to rise substantially thereafter. To accomplish the *Review*'s ambitious goal, CO_2e emissions would need to be progressively cut by $\approx$3% each year, beginning within about a decade. 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?

In this paper I make five basic points about the economics of climate change. 1: The discount rate we choose is all important. 2: We are a lot less sure about core elements of discounting than we commonly admit because critical puzzles and ambiguities are yet unresolved. 3: Structural uncertainty that manifests itself in the thick tails of probability distributions – not risk – is what likely matters most. 4: Gathering information about thick-tailed uncertainties representing rare disasters should be a research priority. 5: The *Stern Review* may be right for the wrong reasons. My overarching theme is that spending money to slow global warming should not be conceptualized primarily as being about optimal consumption smoothing, so much as it should be analyzed more 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.

2 Interest Rates and Long Term Discounting

I think it is fair to characterize the *Stern Review* as being biased in the sense of consistently leaning towards assumptions and formulations that emphasize high damages from climate change and low mitigation costs. But far more crucially, the assumptions that drive the strong conclusions have largely to do with the mundane fact that a very low interest rate is assumed, with which distant-future benefits and costs are then discounted. The eventual upward-sloping "policy ramp" of ever-tighter emissions reductions in the majority of other models (but not beginning now, please) is a familiar consumption-smoothing consequence of discounting: the higher the interest rate the stronger the desire to move towards getting more pleasure now at the expense of postponing even greater pain until later. The *Stern Review* simultaneously raises overall greenhouse-gas reductions and flattens out its consumption-smoothing time profile of pain and pleasure 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 appreciated by economists at large. The insight that the strong conclusions of the *Review* are mainly driven by the low assumed discount rate has been picked up and commented upon already by several insider critics. 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 dynamics for the purposes of analyzing the economic effects of global climate change. An IAM is essentially a model of optimal economic growth with a controllable externality of endogenous greenhouse warming. Behind the *Review* is an IAM called PAGE, on which numbers have been crunched and conclusions have been based. IAMs (including PAGE) have a deterministic optimizing core and treat uncertainty (when it is done formally at all) by taking numerical averages via simulations over imposed probability distributions of alternative parameter values. In Section 4, I will explain that treating risk in this Monte-Carlo way fails to account adequately for the fact that properly integrating out the probability distributions of the structural parameters of unknown probability distributions inevitably causes thick tails in the resulting reduced-form probability distribution, which of necessity must dominate any economic analysis based on expected-utility theory. When uncertainty is central to the model in this sense of probabilities of probabilities causing thick reduced-form tails, investments in slowing climate change then become naturally seen as buying insurance against catastrophic losses with small probabilities that are inherently difficult to estimate. In this situation, the discount-rate issue plays itself out somewhat differently than a debate about the optimal trajectory of consumption smoothing.

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 is the famous Frank Ramsey equation

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

where r is *the* interest rate (more on *the* interest rate later), ρ 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 ρ and η capture two critical aspects of "tastes" while the reduced-form representation of "technology" is the growth rate of consumption g. The important distinction between ρ and r is that ρ is a more-primitive rate of pure time preference that *discounts utility*, while r is the much-more-familiar interest rate used to *discount goods*, which is derived from the more-primitive underlying parameters of tastes and technology. 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 whatever is the ultimate underlying source of g. In Ramsey's time g was conceptualized as coming from capital accumulation, and therefore in long run equilibrium $g \to 0$ and $r \to \rho$. We now know from modern post-Solow growth theory (but Ramsey and Fisher didn't), that in steady-state equilibrium g is essentially the assumed underlying growth rate of labor-augmenting technological progress that, behind the scene, is pushing the entire economy forward (or at least this is so in a world without greenhouse warming). 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): ρ, η, g and r. I will ultimately argue that in a greenhouse-gas world g needs to be seen as a random variable with a climate-change-thickened left tail.

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. Assume that while the absolute economic impacts of global warming about a century from now might be huge, they are still likely to be small relative to the size of the world economy a hundred years hence, which is going to be determined largely by the intervening rate of technological progress, more or less independent of the economic impacts of global warming. For the purposes at hand we are now 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 lies 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 global warming ignores or distorts truly huge chunks of reality. As just one example among many, an important part of the global warming story concerns the big stock-flow lag and enormous built-in inertias from having a long pipeline between greenhouse-gas emissions and temperature change. This built-in inertia is what causes ΔT to continue to rise along the gradualist ever-tighter-emissions majority trajectory to levels well above $\Delta T \approx$2-3°C *after* a century from now. Built-in inertia also fuels the *Stern Review* passion for severe curtailment of greenhouse gas emissions to begin soon because at projected flow rates we will attain a stock of 550 ppm of CO_2e within a half-century and move irreversibly beyond thereafter. 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 – and what it all means. If ever there was a case where boring nerdy details in a simple one-equation economic model actually matter, this is it.

Going right to the target here, my own point-guesstimate of what most economists think are the "correct" parameter values (my ideal thought experiment being the median response from a survey of *AER* or *EJ* readers) would be a "trio of twos": $\rho = 2\%$, $g = 2\%$ (both on an annual basis) and $\eta = 2$. I trust everyone agrees that these numbers at least pass the laugh test. 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 it is about right, so we can temporarily shelve this issue. Plugging these primitives into (1) makes "the" annual interest rate be $r = 6\%$.

Concerning pure time preference, *Stern* follows a decidedly-minority paternalistic view (which, however, includes several distinguished economists) that for social discounting selects the lowest conceivable value $\rho \approx 0$ according to the *a priori* philosophical-ethical principle of treating all generations equally – irrespective of preferences for present over future utility that people seem to exhibit in their actual saving and investment behavior. Likewise, *Stern* chooses as a coefficient of relative risk aversion the value $\eta = 1$ that is the lowest lower bound of just about any economist's best-guess range. The preferred point estimates of the *Stern Review* are $\rho = 0.1\%$ p.a., $g = 1.3\%$ p.a., $\eta = 1$, which makes "the" annual interest rate from (1) be $r = 1.4\%$. The present discounted value of a given global-warming loss from a century hence at the non-*Stern* annual interest rate of $r = 6\%$ is *one hundreth* of the present discounted value of the same loss at *Stern*'s annual interest rate of $r = 1.4\%$. The disagreement over what interest rate to use for discounting is equivalent 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 1% to 4%. Consistent with its philosophy of picking the gloomiest possibilities as base case scenarios in a heuristic attempt to include fearsome extreme possibilities – because "when we try to take due account of the upside risks and uncertainties, the probability-weighted costs look very large" – the *Stern Review* selects $\frac{D}{Y} > 5\%$ (and actually goes well beyond 5% in a spiritual-intuitive, if not numerical-analytical, sense by making scattered literary and numerical allusion to the dark possibilities lurking in the tails of the distribution of possible outcomes). *Stern* also estimates the annual costs of its ambitious abatement strategy as being equivalent to about 1% of GDP (which seems somewhat on the low side, but that is not so relevant here). The question for the *Stern Review* analysis then effectively becomes: is it worthwhile to sacrifice 1% of GDP now to save 5% 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), $r - g = (\eta - 1)g + \rho$, so that by picking the extreme values $\eta = 1$, $\rho = 0.1\%$, the *Stern Review* guarantees that the difference $r - g$ is *always* the miniscule value $\rho = 0.1\%$, 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. (One wonders why the *Stern Review* did not make life easier by just rounding down a mere tenth of a percent by assuming $\rho = 0$, which along with $\eta = 1$ would then make cost-benefit analysis really simple here because a fixed fraction of GDP would then be worth the same fixed fraction of GDP at *any* future time and for *any* projected growth rate.) With the *Stern Review* 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 $g = 2\%$, $r = 6\%$ make $\frac{B}{C} = \frac{1}{10}$ – a clear reject. This simple 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.

The *Stern Review* taste-parameter values are much too low to be consistent with observed savings behavior. For a neoclassical optimal growth model with labor-augmenting technological progress growing at rate g, if s is the savings rate out of net income then in steady-state equilibrium $s = \frac{\dot{K}}{Y} = \frac{gK}{Y} = \frac{rK}{Y}\frac{g}{r}$, or

$$s = \frac{\alpha g}{\rho + \eta g}, \tag{2}$$

where α is the share of capital. Suppose $\alpha = \frac{1}{3}$. Plugging the *Stern Review* parameter values $\rho = 0.1\%$, $g = 1.3\%$, $\eta = 1$ into formula (2) gives $s = 31\%$, which is distressingly higher than what is observed. Using the alternative trio of twos $\rho = 2\%$, $g = 2\%$, $\eta = 2$ in formula (2) gives $s = 11\%$, which is much closer to what might pass as reality for savings rates. This kind of savings-rate calibration exercise lends credence to numbers like the "trio of twos" and tends to make most mainstream economists have doubts about *Stern*'s assumed parameter values. Several critics of the *Stern* taste parameters inappropriately employ the pre-Solow original 1928 Ramsey model (in which capital is the only source of growth and $g \to 0$) to argue that since $s = 1/\eta$ when $\rho = 0$ along the Ramsey optimal-growth trajectory, there is a *reductio ad absurdum* that on the surface seems devastating for the *Review* taste parameters – because then $\eta \approx 1$ implies $s \approx 100\%$. However, this overkill formulation itself enormously contradicts stylized facts about economic growth by omitting the prime driver: labor-augmenting technological progress. As will later be elaborated, there are other empirical calibrations of taste parameters to stylized economic facts, which arguably are far more directly relevant to discounting investment possibilities than observed savings behavior and for which *both* sets of taste parameters (*Stern* and non-*Stern*) fail miserably by several orders of magnitude.

3 Ambiguities and Puzzles 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 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 $r = 6\%$ and a $\frac{1}{2}$ chance of $r = 1.4\%$ are not at all the same thing as splitting the difference by selecting the average $r = 3.7\%$. 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 $.5e^{-6} + .5e^{-1.4}$ a century hence, which, when you do the math, is equivalent to an *effective* interest rate of $r = 2\%$. 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 $r = 1.4\%$ and the more conventional value of $r = 6\%$, but with the above numbers it is a lot closer to the *Stern* value and is not anywhere near the arithmetic average of $r = 3.7\%$. 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 uphill profile 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 world of the Ramsey deterministic formula (1), there is no distinction among rates of return on various assets and r is just *the* economy-wide rate of return on capital or, more succinctly, *the* interest rate. In 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 economy-wide rate of return applicable to investments that have payoff characteristics parallel to the economy itself and a riskfree 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 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

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

where r^f in equation (4) denotes the *riskfree* 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 a claim on the consumption dividend produced by the macroeconomy itself. 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. Abstracting away from taxes, liquidity, leverage, and so forth, a famous formula from the famous Lucas-fruit-tree general equilibrium model shows that g being i.i.d. $N(\mu, \sigma^2)$ makes the equity risk premium over the safe rate be

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

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. Combining (5) with (4) gives the average return on equity as

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

Extending the previous "trio of twos" parameter values to a not-implausible knee-jerk "quartet of twos" $\rho = 2\%$, $\eta = 2$, $E[g] = 2\%$, $\sigma[g] = 2\%$ makes very little difference on the riskfree rate because now $r^f = 5.9\%$ in (4) instead of the previous value for "the" interest rate of $r = 6\%$ in (1). The corresponding equity premium from (5) is $\overline{r}^e - r^f = 0.1\%$ and the average return on equity from (6) is $\overline{r}^e = 6\%$. The actual empirical numbers are closer to $r^f \approx 1\%$, $\overline{r}^e - r^f \approx 6\%$, $\overline{r}^e \approx 7\%$. (The calibration $r^f \approx 1\%$ refers to short-term treasury bills, while $\overline{r}^e \approx 7\%$ 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 riskfree rate and the equity premium – thereby giving rise to the notorious "riskfree 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 numbers temporarily aside in favor of first asking a fundamental pre-numerical question: *in principle* (leaving aside their correct numerical values) should we be using the riskfree rate or the risky economy-wide 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_2 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 payoffs R^e and R^f, not r^e and r^f). It then follows from essentially the same considerations as went into deriving formula (5) that the relevant interest rate for discounting costs and benefits at time t here is

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

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 formally lies behind the *Review* conclusions, by a particular sub-aggregator equation of the multiplicative form

$$D(t) \;=\; Y^*(t) - Y(t) \;=\; f(\Delta T(t))\, Y^*(t), \tag{8}$$

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 ("business as usual" with $\Delta T \approx 3°$C) be somewhere between approximately 0% and approximately 5% (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% (representing economy-wide average returns) or the value of 6% predicted by formula (5) from my non-*Stern* "quartet of twos" parameter values – but the discrepancy between 6% and 7% is insignificant for my 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 consumption smoothing and bring the results back to the much more moderate take-it-more-slowly majority-tilt time profile advocated by the mainstream critics of *Stern*.

This important debate about what interest rate to use for discounting costs and benefits of mitigating greenhouse warming mirrors a previous debate on the same subject over a decade ago between Cline and Nordhaus, two early pioneers of modeling the economic effects of climate change. Like *Stern*, the 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 $\rho = 0\%$, $\eta = 1.5$, $g = 1\%$, which combined to make *the* interest rate be $r = 1.5\%$ per year. Like *Stern*, the strong activist conclusions of Cline's work a dozen years earlier also traced back to the very low discount rate being used. Furthermore, Cline and *Stern* are soul-mates in their *cri de coeur* justifying $\rho \approx 0$ 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. Inferring society's revealed preference value of

ρ is not an easy task in any event, but at least a good faith effort to do so would go 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 (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 $r = 1.5\%$ is calibration-consistent with the real return on U.S. Treasury bills historically being about 1% 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 IAM's has long set a standard in this arena, argued forcefully over a decade ago that the riskfree interest rate should not be used for discounting costs and benefits of climate change. In this argument Nordhaus was following 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."

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 circuitous route of being multiplicative in GDP according to the particular sub-aggregator function (8). It was more due to an historical accident of stumbling upon a convenient analytical form whose parameters could be 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, more relevantly here, perhaps entering the utility function as a direct argument – 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 "outdoor" aspects like agriculture, coastal 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 determined 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. The relevant share of this "outdoor" subset of the economy in investment-beta calculations might be disproportionately large because it is disproportionately-largely impacted by greenhouse warming. What then happens to the discount rate for climate naturally depends on the actual value of β that is assumed.

If $\beta = 0$ in (7), then $r(t) = r^f = 1\%$. If $\beta = 1$ in (7), then $r(t) = \overline{r}^e = 7\%$. 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 a correlation coefficient of, say, $\beta \approx .5$. With $\beta = .5$ in (7), the relevant interest rate for a century from now becomes $r(100) = 1.7\%$, which is close to Stern's $r = 1.4\%$ or Cline's earlier $r = 1.5\%$. 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 $\beta = .5$. Rather, the moral of this story is that the nature of the impacts of climate change determine whether we should end up closer to the risk-free rate or the economy-wide 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 riskfree interest rate, which is close to the *Stern* interest rate, 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 – and from the same underlying economic logic that makes compound interest a powerful force in the first place.

Next, suppose we try to repeat the above numerical exercise but in place of the empirical values $r^f = 1\%$, $\overline{r}^e = 7\%$, we use the values predicted by the theoretical formulas via assuming the "quartet of twos" parameter values, which then implies $r^f = 5.9\%$ from (4) and $\overline{r}^e = 6\%$ from (6). Because the equity premium predicted from (5) is a miniscule 0.1%, there is essentially no difference in this case between $r^f = 5.9\%$ and $\overline{r}^e = 6\%$. The relevant discounting rate $r(t)$ from (7) then lies between 5.9% and 6% independent of the assumed value of β. When $5.9\% \leq r(t) \leq 6\%$, 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 which interest rate to use for discounting costs and benefits of climate change thus becomes tied up with the interpretation of the equity-premium and riskfree-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 one direction, then r is sensitive to β and $r \approx 1.7\%$ for $\beta = .5$. If we use numbers that resolve the asset-pricing puzzles in the other direction, then $r \approx 6\%$ independent of β. And, to whip a horse long dead, it makes a huge difference to the economics of climate change whether $r \approx 1.7\%$ or $r \approx 6\%$.

Critics of the *Stern Review* are fond of pointing out that $\rho \approx 0$, $\eta \approx 1$ is inconsistent with observed economic behavior, especially savings behavior. While this is true, 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 numbers that theory predicts we should be using for discounting and the actual discount-rate numbers that are 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 ρ and η that we are commonly using (here $\rho = 2\%$ p.a. and $\eta = 2$) are wrong. If we treat (4) and (5) as two equations in two unknowns (ρ and η), we can then invert the two equations to back out the hypothetical values $\widehat{\rho}$ and $\widehat{\eta}$ that would "explain" the stylized-fact empirical observation that $r^f \approx 1\%$ and $\overline{r}^e - r^f \approx 6\%$. When this is done (for $\mu = 2\%$, $\sigma = 2\%$), it produces the mega-puzzle that the estimated rate of pure time preference is $\widehat{\rho} \approx 151\%$ per year and the coefficient of relative risk aversion is $\widehat{\eta} \approx 150$. One does not know whether to laugh or to cry at the prospect of what the *Stern Review* IAM might end up recommending as the preferred policy for climate change in its number-crunching simulations if the parameter values $\widehat{\rho} = 151\%$, $\widehat{\eta} = 150$ were plugged into PAGE. So much for the fantasy that values of the taste parameters ρ and η should be chosen to be consistent with the revealed-preference 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 the optimal growth model. 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 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 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 a 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 greatly 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 finite, are naturally limited in coverage. Besides, even if we had an infinite time series of past observations, they are of limited relevance in an evolving world where the environment is changing and the 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. Such an effect could end up significantly lowering r^f relative to the observed past average of realized r^e.

There is little doubt that the worst-case scenarios of global warming are genuinely frightening. The *Stern Review* goes over several of these poorly-understood threshold-crossing highly-unlikely disasters associated with abrupt large-scale irreversible changes in the climate system: sudden collapse of the Greenland and West Antarctica ice sheets, weakening or even reversal of thermohaline circulations that might strongly affect the Gulf Stream and European climate, runaway amplification of global warming due to the many potential reinforcing feedbacks (including, but not limited to, loss of polar albedo and rapid releases of methane from arctic permafrost). More gradual but still very serious and highly uncertain are: sea-level dynamics, drowned coastlines of unknown magnitude, extreme weather patterns, flood risks, ecosystem destruction, mass species extinctions, tropical-crop failures, 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. Like everything else, ΔT has a probability distribution. The just-released executive summary of the *Fourth Assessment Report of the IPCC* predicts for one hundred years from now a mean temperature change from six equally-plausible scenario groups of $\Delta T \approx 2.8°C$ with standard deviation $\approx$1.6°C (Table SPM-2). This means the probability that $\Delta T > 4.5°C$ is approximately 15% and the probability that $\Delta T > 6°C$ is about 2%. Societies and ecosystems whose average temperature has changed in the course of a century by $\Delta T > 6°C$ (for U.S. readers: $\Delta 6°C \approx \Delta 11°F$) are located in the *terra incognita* of what any honest economic modeler would have to admit is a planet Earth reconfigured as science fiction.

The idea behind analyzing projects by reducing future costs and benefits to present discounted values is that society has alternative investment opportunities, whose rate of return is the discount rate. These alternative investments represent capital accumulation throughout the rest of the economy that would compensate us for the economic losses suffered from climate change. Mundane examples of alternatives to CO_2e mitigation from middle-of-the-probability-distribution mild warming might include accumulating air conditioners 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 non-market enjoyment of the natural environment – like habitats, ecosystems, and species – then it is difficult to imagine what are the compensating alternative investments. With 2% IPCC-4 probability, we will "consume" a *terra incognita* natural landscape within a hundred years whose mass species extinctions will have been triggered by a geologically-instantaneous temperature change that is larger than what separates us now from past ice ages.

In the rest of this paper 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 costs of low-g rare disasters from high-ΔT scenarios more properly constitute uncertainty in the sense of Knight or Keynes than risk because it is so very difficult to make probability estimates of high-ΔT catastrophic damages – due, ultimately, to the underlying sampling-theory principle that the rarer is an event the more unsure is our estimate of its probability. 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 probabilities in the tails. The net result is thicker left tails for growth rates 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 supposed to be 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 this purpose – I think that contemporary economic practise goes too far and leads to a mindset that all-too-easily identifies subjective probabilities with 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 near the middle because, from previous experience, past observations, plausible extrapolations, and, perhaps, 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 towards probabilities in the tails of the g distribution, however, we are increasingly moving into the unknown territory of subjective uncertainty where the probability distribution of our estimates of the probabilities 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 of the probability distribution of g (after integrating out the probabilities of probabilities) – which is relevant for the economics of climate change because 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 on how bad a catastrophe global warming might induce and with what probability, but also on how imprecise are our probability estimates of the probabilities of those bad catastrophes. Mitigating the future consequences of greenhouse warming does not just shift the center of the distribution of g to the right but, perhaps 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 being uncertain about uncertainties 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 $\mu = 2\%$ and $\sigma = 2\%$. 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 riskfree 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 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 $\widehat{\mu} = 2\%$ and $\widehat{\sigma} = 2\%$ from this regression just so happen to be the very same numbers as the 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 unbounded. The bombshell fact that, as soon as we admit that we don't know the underlying stochastic structure and parameters must be estimated, $EMU = +\infty$ changes the rules of the game. This fact represents a mathematically-generic result, which is 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 that are not easily brushed aside, and not least of all for cost-benefit analysis of climate change.

There is a general point here and a particular application to greenhouse warming. The general point is that from limited empirical data 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. (This mechanism explains the asset-return puzzles for reasonable values of ρ and η as being due to a fear of relatively rare tail disasters that is theoretically difficult or impossible to eliminate when the underlying structure remains at all uncertain.) Applied to climate change, structural uncertainties about what global warming will do to the left tail of the g-distribution are quite naturally modeled as if caused by limited experience, which makes it all but inevitable that this aspect is critically important because it effortlessly drives the outcome of any cost-benefit calculation. Therefore, to ignore or suppress the significance of rare tail disasters in an application of expected utility theory like cost-benefit analysis of climate change, where there is so obviously limited data and limited information about the evolving structure of the global-warming transformation that is unfolding around us, 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 any model, when pure economic theory proves 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 upon whomever wants to model optimal-expected-utility growth under endogenous greenhouse warming *without* having uncertainty tending to matter more than risk. Such a center-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, does *not* play a decisive role in the analysis.

5 Climate Uncertainty and the Value of Information

Because the *Stern Review* is imbued with the laudable moral imperative not to subject future generations to the tribulations of global warming, it does not shy away from emphasizing (at least verbally) the possibilities of rare high-ΔT, low-g left-tail 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 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 $\rho \approx 0$ and of $\eta \approx 1$ (and, secondarily, such high values of $\frac{D}{Y} > 5\%$) as will operate through the back door of conventional economic analysis to weight distant-future damages high enough 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 $\rho \approx 0$. The "ethical judgements on the distribution of income" seem like *Stern Review* code language for subliminally picking values of η to support the activist conclusions that it wants subconsciously to reverse engineer. Here the *Review* really seems like it is 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 large) calls for urgent action now to avoid future massive spatial redistribution of income from the poor to the rich. Simultaneously, the *Review* wants to further exacerbate distributional inequities by redistributing income *across time* from the relatively-poor present to the relatively-rich future a century or so from now (when the standard of living will likely be ten times higher) via choosing the lowest imaginable value $\eta = 1$ (along with $\rho \approx 0$) that might formally be used to reverse-engineer the really low discount rates it needs to prop up its technical case for immediate urgent action.

I think that rather than trying to go though the back door with unreasonably low values of ρ and η, (or, secondarily, "averaging across possibilities" by heuristically making business-as-usual $\frac{D}{Y} > 5\%$ instead of some smaller more-plausible point estimate), it is much better to go directly through the front door with the legitimate issue that there is a chance, whose subjective probability is small but diffuse (thereby resulting in a dangerously-thickened left tail of growth rates), that global warming may cause disastrous environmental changes. If one accepts this assessment, then many hard questions need to be asked. What are early-warning signs of impending sudden environmental disasters like melting ice sheets or thermohaline inversions? How much would it cost to put sensors in place that might detect 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 an early-warning signal? Would last-ditch emergency measures to ward off disaster by reversing the worst consequences of global warming come 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 artificial dust or sulfate aerosols.) Could such last-ditch measures be made reversible by building in decay mechanisms (as with sulfate aerosols, while we then use the new information to really *really* undertake draconian measures to cut greenhouse gas emissions drastically)? Can the public accept the politically-incorrect idea, which is a third rail few policy-makers dare to touch, that in the extremely unlikely event of an truly-extraordinary unfolding disaster it might be a good backup plan to purposely tinker with spaceship Earth yet further? (This time our purposeful geo-engineering would be to offset the unintended bad consequences of our previous tinkering with the planet via an industrial revolution that turned out to have delivered near to medium-term prosperity at the long-term cost of having burned so much fossil fuel that with hindsight we inadvertently created an environmental catastrophe.)

I trust that some readers can 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 from now to thinking more about what might be in the tails with $\Delta T > 6$°C, which is just two IPCC-4 standard deviations away (i.e., probability $\approx$2%) and which 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. ("Nondogmatic" in Bayesian parlance just means that no event is ruled out a priori by being assigned zero probability in the prior distribution.) A responsible policy approach neither dismisses the horror stories just because they are a couple of 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 standard 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 worst-case trajectory (and knows how much early-warning detection sensors cost to install), confronts the possible options of undertaking politically-incorrect emergency measures if a worst-case nightmare trajectory happens to materialize, and generally attempts be constructive by having some semblance of a game plan for dealing realistically with what might possibly be coming down the road. The point is to supplement traditional mid-probability-distribution analysis of, and action on, climate change by putting serious research dollars into early detection of rare disasters and 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 undoing greenhouse warming anyway – so we should stop stalling and now begin making serious down payments on catastrophe insurance. But these are conclusions we need to reach empirically, rather than prejudging them initially.

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 bequeath to our great-great-grandchildren 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, the *Review* does not follow through formally on this really unsettling part of the global warming equation – which a generous interpretation of its not-great economic analysis might say is the underlying motivation for its overall alarmist tone – except indirectly, by choosing $\rho \approx 0$, $\eta \approx 1$, $\frac{D}{Y} \approx 5\%$ in order to reverse-engineer the drastic slowing measures that it intuitively and subconsciously 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 meaning of *Stern vs. Critics* may be about *left-tail 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 – 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 politician planning to remain in office anywhere wants to hear, that the world needs desperately to start confronting the reality that burning carbon has a significant externality cost that should be taken into account by being charged full-freight for doing it. (This should have been, but of course was not, the most central "inconvenient truth" of all in Al Gore's tale about inconvenient climate-change truths.) As the *Review* puts it, "establishing a carbon price, through tax, trading, or regulation, is an essential foundation for climate-change policy." One can only wish that U.S. political leaders might have the wisdom to understand and the courage to act upon the breathtakingly-simple relatively-market-friendly idea that the right carbon tax could do *much* more to unleash the decentralized power of greedy, self seeking, capitalistic American inventive genius on the problem of developing commercially-feasible carbon-avoiding alternative technologies than all of the command-and-control schemes and patchwork subsidies making the rounds in Washington these days. As I have made clear here, a generous interpretation might also credit the *Stern Review* with intuiting the greater significance of insuring against catastrophic uncertainty than of consumption smoothing for the climate problem, even if this intuition remains subliminal and does not formally enter the analysis through the front door.

To be honest about 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 $r \approx 1.4\%$ via the indirect route of picking parameter values $\rho \approx 0$ and $\eta \approx 1$ that are more like theoretically-reasoned extreme lower bounds than empirically-plausible estimates of representative tastes. In this sense, it must be said staightforwardly that the subconsiously-reverse-engineered output of *Stern*'s PAGE simulations and the goal-oriented formal economic analysis of the *Review* are not worth a great deal. 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 $r \approx 6\%$ to much lower values of perhaps $r \approx$ 2-3%, 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 its mainstream 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 that are capable of driving expected-utility outcomes.

On the political side of the *Stern Review*, my most charitable interpretation of its urgent tone is that the report is an essay in persuasion that is more about gut instincts regarding the horrors of uncertain rare disasters whose probabilities we do not know than it is about economic analysis as that term is conventionally understood. Although it is difficult enough to analyze people's motives, much less the motives of a 600-page document, I can't help but think after reading it that the strong tone of morality and alarm is mostly reflecting a fear of what is potentially out there with greenhouse warming in (using ponderous terminology here to make sure the thought is exact) "the inherently-thick 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-thick 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* was more wrong or more right, and, if it was more right, whether as pure economic analysis it was right for the right reasons or it was right for the wrong reasons.