

RISK AND RETURN IN ENVIRONMENTAL ECONOMICS*

by

Robert S. Pindyck
Massachusetts Institute of Technology
Cambridge, MA 02142

This draft: July 23, 2012

Abstract: I examine the risk/return tradeoff for environmental investments, and its implications for policy choice. Consider a policy to reduce carbon emissions. To what extent does the value of such a policy depend on the expected future damages from global warming versus uncertainty over those damages, i.e., on the expected benefits from the policy versus their riskiness? And to what extent should the policy objective be a reduction in the expected temperature increase versus a reduction in risk? Using a simple model of a stock externality (e.g., temperature) that evolves stochastically, I examine the “willingness to pay” (WTP) for alternative policies that would reduce the expected damages under “business as usual” (BAU) versus the variance of those damages. I also show how one can compute “iso-WTP” curves (social indifference curves) for combinations of risk and expected returns as policy objectives. Given cost estimates for reducing risk and increasing expected returns, one can compute the optimal risk-return mix for a policy, and the policy’s social surplus. I illustrate these results by calibrating the model to data for global warming.

JEL Classification Numbers: Q5; Q54, D81

Keywords: Environmental policy, climate change, economic growth, risk, uncertainty, risk-return tradeoff, willingness to pay, policy objectives.

*My thanks to Stacie Cho, Lorna Omondi, and Andrew Yoon for their excellent research assistance, and to John Cox and seminar participants at M.I.T., University of Quebec at Montreal, and Tel-Aviv University for helpful comments and suggestions.

1 Introduction.

The adoption of an environmental policy usually imposes costs on society, but is expected to yield a social return in the form of a stream of benefits, e.g., health benefits from less pollution. Those benefits, however, usually occur in the future and may be highly uncertain, so that (as with other private or public investments) the actual return is uncertain. This is especially true for environmental policies involving stock externalities, such as increases in greenhouse gas (GHG) concentrations, the acidification of lakes and oceans, and the accumulation of toxic waste. All of these problems involve long time horizons and uncertain impacts. This paper examines the risk/return tradeoff for environmental investments and the implications of that tradeoff for policy choice. I utilize a very simple and stylized model which focuses on climate change, because of the long time horizon and considerable uncertainty involved. However, the framework developed here could easily be applied to other environmental (and non-environmental) policy problems, such as the management of toxic (including nuclear) waste, and the development and protection of water resources.¹

Consider a policy designed to reduce the accumulation of atmospheric greenhouse gases, e.g., via the imposition of a carbon tax. That policy would have some cost to society, but by reducing GHG emissions it would reduce the extent of warming by some uncertain amount, and thus would yield uncertain future benefits. The uncertainty arises because we have limited knowledge of the relationship between GHG concentrations and temperature, and we know very little about how higher temperatures would affect GDP growth or other indicia of welfare. Thus we could ask two questions. First, how large are the expected social benefits from a given policy relative to the risk involved? Second, to what extent is the social value of the policy (measured by society's willingness to pay for it) driven by expected benefits from the policy versus uncertainty over those benefits.

The answers to these questions can guide policy design. Alternative policies might have different impacts on the expected change in temperature over the next century versus the

¹For a general discussion of the nature of the uncertainties involved in environmental policies, whether or not stock externalities are involved, see Pindyck (2007).

variance of that change. For example, rather than (or in addition to) GHG emissions abatement, money might be spent to improve our knowledge of how GHG emissions affect temperature, and how changes in temperature affect economic output. Therefore a purely research-focused policy might do little or nothing to reduce the expected change in temperature (and hence the expected impact on GDP), but might reduce the variance of our estimates of future temperature and GDP. Likewise, investments in adaptation to climate change (e.g., the development of hardier hybrid crops) would reduce the right tail of the distribution of future damages, reducing both the variance and expectation of damages.

As a general matter, there is value in reducing the expected damages from GHG emissions *and* in reducing the variance of those damages. At issue is the tradeoff between the two. To what extent should policy be aimed at reducing expected damages versus the variance of those damages?

The framework I use to address this question is “willingness to pay” (WTP). Consider a policy that would reduce expected future temperature increases, and possibly reduce the uncertainty over those temperature increases. The willingness to pay for the policy is the *maximum* percentage reduction in current and future consumption that society should be willing to give up in order to achieve those reductions in the expectation and/or variance of future temperature increases.² This does not mean it will be possible to achieve these reductions given the revenues from this WTP; doing so may be more costly, or less costly. WTP relates only to the *demand side* of policy, i.e., society’s reservation price for the achievement of particular policy objectives.

In Pindyck (2011a, 2012), I calculated WTP for GHG abatement policies that would limit future increases in temperature to some specified amount (e.g., 3°C), based on a probability distribution for the rate of warming absent the policy, i.e., under “business as usual” (BAU). Of course WTP depends on the probability distribution, and in particular depends on both the *expected* increase in temperature under BAU and the *variance* of the increase in temperature. I calculated the trade-off between the expected value of a temperature increase

²I used a related measure of WTP to examine climate change policies under uncertainty; see Pindyck (2011a, 2012). WTP has also been used in the context of climate change by Heal and Kriström (2002).

versus its standard deviation under BAU as determinants of WTP. That trade-off is the marginal rate of substitution along an “iso-WTP curve,” i.e., the locus of combinations of the expected value versus the standard deviation of the change in temperature under BAU such that the WTP (e.g., to prevent any temperature increase) is constant. The results showed that uncertainty over temperature change can be a stronger driver of WTP than expectations, and thus should be a major focus of climate change policy.

That earlier work focused on the *starting conditions* for policy, i.e., the relative contributions to WTP of the expected value versus the variance of the change in temperature under BAU. Here I turn to the *objectives* of policy, and examine the trade-off between reducing the expected rate of the increase in temperature versus reducing the variance of future increases in temperature. In particular, what combinations of drift reduction and variance reduction yield the same WTP? Those combinations trace out a social indifference curve — combinations of changes in drift and variance that are welfare-equivalent.

Given estimates of the costs of drift reduction and variance reduction, we can take this a step further and determine whether the cost of a particular combination of drift and variance reduction is less or greater than the WTP for that combination. If the cost is less than the WTP, the policy yields a positive social surplus. If the cost exceeds the NPV, the policy is not economically feasible. Finally, assuming the costs are convex (as we would expect), we can determine the optimal risk-return tradeoff for policy design.

In the next section I lay out the model that I use to address these questions. Despite its simplicity, the model captures the key features of environmental policy risk and expected returns that are of interest. In Section 3, I use the model to analyze risk and return for a policy that reduces the expected rate of temperature increase by some incremental amount. Section 4 addresses policies to change the expected rate of temperature increases and/or the variance of future temperature increases, and shows how we can calculate the willingness to pay for such policies. In Section 5, I calibrate this model to information on global warming and its impact, and then use it to address risk-return tradeoffs.

2 A Simple Model.

I use a simple model in which an environmental stock variable, in this case temperature change, follows an arithmetic Brownian motion (ABM) with positive drift and some volatility.³ The impact on consumption is through the real growth rate of GDP, which is linearly related to the change in temperature. Welfare is measured using CRRA utility. Various expectations and variances of utility can be found analytically, making the model very tractable.

2.1 Temperature and GDP Growth.

Let X_t denote the anthropomorphic increase in temperature from its current level. I assume that under “business as usual,” X_t follows an arithmetic Brownian motion (ABM) with drift and volatility α_X and σ_X respectively. I also assume that the impact of an increase in X_t is to reduce the real rate of growth of consumption, g_t . This differs from the damage function specifications used in most integrated assessment models of climate change, where X_t is assumed to reduce consumption directly.⁴ Theoretical arguments and empirical evidence support the view that higher temperatures (and environmental damage in general) should affect the growth rate of consumption rather than its level, as in the model I use here.⁵

³In the context of global warming, the actual stock variable is the atmospheric GHG concentration, which in turn drives temperature change, but with a lag. I ignore the lag, and thus can treat temperature change itself as the stock variable. In the context of toxic waste, the stock variable could be the quantity or concentration of waste material over some geographic area.

⁴Most economic studies of climate change relate the temperature increase T to GDP through a “loss function” $L(T)$, with $L(0) = 1$ and $L' < 0$, so GDP at a horizon H is $L(T_H)GDP_H$, where GDP_H is but-for GDP with no warming. These studies — see, e.g., Nordhaus (2008) — typically use an inverse-quadratic or exponential-quadratic function. The loss function $L(T)$ implies that if temperatures rise but later fall, GDP could return to its but-for path with no permanent loss.

⁵There are several reasons to expect higher temperatures to affect the growth rate of GDP as opposed to the level. First, some effects of warming will be permanent; e.g., destruction of ecosystems, extinction of species, and deaths from weather extremes. A growth rate effect allows warming to have a permanent impact. Second, there is theoretical support for a growth rate effect. The resources needed to counter the impact of warming will reduce those available for R&D and capital investment, reducing growth. Adaptation to rising temperatures is equivalent to the cost of increasingly strict emission standards, which, as Stokey (1998) has shown with an endogenous growth model, reduces the rate of return on capital and lowers the growth rate. (Also see the related study by Brock and Taylor (2010).) Finally, there is empirical support for a growth rate effect. Using historical data on temperatures and precipitation over 50 years for a panel of 136 countries, Dell, Jones, and Olken (2009, 2012) have shown that higher temperatures reduce GDP growth rates but not levels. Also, using economic and financial market data for a panel of 147 countries over the

The process for X_t is thus:

$$dX_t = \alpha_X dt + \sigma_X dz , \quad (1)$$

where dz is the increment of a Weiner process, and $X_0 = 0$. Note that eqn. (1) is essentially a vehicle for describing the expectation and variance of the temperature change from $t = 0$ to any future point in time T . In particular, $\mathcal{E}_0(X_T) = \alpha_X T$ and $\text{Var}(X_T) = \sigma_X^2 T$.

The growth rate of consumption is given by the following linear relationship:

$$g_t = g_0 - \gamma X_t , \quad (2)$$

where g_0 is the growth rate absent any change in X_t . Thus g_t also follows an ABM:

$$dg = -\gamma\alpha_X dt - \gamma\sigma_X dz \equiv -\alpha dt - \sigma dz . \quad (3)$$

Eqn. (2) implies that at any time s ,

$$g(s) = g_0 - \alpha s - \sigma \int_0^s dz = g_0 - \alpha s - \sigma z(s) ,$$

so consumption at a future time t can be written as:

$$C_t = C_0 e^{\int_0^t g(s) ds} = C_0 e^{g_0 t - \frac{1}{2} \alpha t^2 - \sigma \int_0^t z(s) ds} , \quad (4)$$

To measure welfare, I use a constant relative risk aversion (CRRA) social utility function:

$$U(C_t) = C_t^{1-\eta} / (1-\eta) , \quad (5)$$

where η is the index of relative risk aversion (and $1/\eta$ is the elasticity of intertemporal substitution). At time 0, welfare (under business as usual) is then given by:

$$W_0 = \frac{1}{1-\eta} \mathcal{E}_0 \int_0^\infty C_t^{1-\eta} e^{-\delta(t-s)} dt \quad (6)$$

where δ is the rate of time preference, i.e., the rate at which utility is discounted.

period 1950 to 2007, Bansal and Ochoa (2011a, b) show that increases in temperature have a strong negative impact on economic growth. For an analysis of the policy implications of direct versus growth rate impact, see Pindyck (2011a).

2.2 Willingness to Pay.

Suppose that with some expenditure over time, society could reduce α_X and/or σ_X , i.e., reduce the expected rate of temperature change and/or the volatility of temperature change (and hence the variance of the change in temperature at any arbitrary date in the future). Let α_X and σ_X be the drift and volatility of X_t under BAU, and α'_X and σ'_X be the corresponding drift and volatility under a policy that has a permanent cost to society of w percent of consumption. Here $\alpha'_X \leq \alpha_X$ and $\sigma'_X \leq \sigma_X$. Willingness to pay is the maximum value of w that society would accept to reduce α_X to α'_X and reduce σ_X to σ'_X .

While a policy that reduces α_X might seem straightforward, how as a practical matter could a policy be designed to reduce σ_X , i.e., to reduce the variance of temperature change T years from now? Investment in research is the most obvious example. As I discuss later in this paper, in the context of climate change research can be expected to provide better estimates of the rate of temperature increase under, say, BAU, and also better estimates of the impact of higher temperatures. Another example is investments in adaptation to higher temperatures. An extreme case of this is the development of geoengineering technologies that could be used to limit or even reverse temperature increases.⁶

Note that WTP is society's reservation price for achieving the policy objective $(\alpha_X, \sigma_X) \rightarrow (\alpha'_X, \sigma'_X)$. However, it is a reservation price for a particular type of policy and a particular form of pricing: At time $t = 0$ the policy is adopted (with no option to wait for more information), and the drift and volatility of X_t are immediately changed to their new values. The payment flow, however, starts at $t = 0$ and continues forever, i.e., takes the form of a permanent reduction of consumption of w percent (yielding a permanent flow of revenue to pay for the policy objective). One might naturally argue that this is unrealistic. First, we would expect that changing α_X and/or σ_X should take time, and perhaps considerable time. Second, there is little reason to expect that the flow cost of achieving this policy objective will be a fixed percentage of GNP or consumption. Thus WTP might instead be defined as

⁶This might be done by seeding the atmosphere with sulfur particles. For a discussion of this and other geoengineering approaches to climate change, see Barrett (2007).

the willingness to give up a particular time-varying percentage of consumption in return for a shift in α_X and/or σ_X that occurs over some period of time. However, the simplification that I use (α_X and/or σ_X change instantly and the flow cost is a constant percentage of consumption) adds clarity to the basic results.

Welfare under the policy is given by:

$$W_1(\alpha'_X, \sigma'_X) = \frac{(1-w)^{1-\eta}}{1-\eta} \mathcal{E}'_0 \int_0^\infty C_t^{1-\eta} e^{-\delta t} dt, \quad (7)$$

where \mathcal{E}'_0 denotes the expectation at $t = 0$ when the drift and volatility of X_t are α'_X and σ'_X . Under BAU (i.e., no policy), welfare is:

$$W_2 = \frac{1}{1-\eta} \mathcal{E}_0 \int_0^\infty C_t^{1-\eta} e^{-\delta t} dt, \quad (8)$$

where \mathcal{E}_0 is the expectation under the original drift and volatility, α_X and σ_X . Then WTP is the value w^* that equates W_1 and W_2 . I treat GDP and consumption as interchangeable, and assume that all losses from higher temperatures, including health effects and ecosystem damage, can be monetized and included in GDP.

From eqns. (2) and (3), the drift and volatility of temperature correspond to a drift and volatility of the real growth rate g_t , i.e., $\alpha = \gamma\alpha_X$, $\sigma = \gamma\sigma_X$, $\alpha' = \gamma\alpha'_X$, and $\sigma' = \gamma\sigma'_X$. In what follows I will refer to changes in α and σ rather than α_X and σ_X .

Suppose that under BAU, the drift and volatility of g_t are α_0 and σ_0 respectively. If we specify a target drift and volatility, α' and σ' , we can calculate the WTP for a policy that would achieve this target. Suppose that WTP is w_1^* . We are interested, however, in the tradeoff between drift reduction and variance reduction as targets of policy. That is, we want to know what combinations of drift reduction and variance reduction will yield the same WTP, w_1^* , assuming we start at α_0 and σ_0 . We can compute the locus of such combinations and thereby obtain an “iso-WTP curve,” i.e., combinations of α and σ which, starting at α_0 and σ_0 , all have the same WTP, and thus are welfare-equivalent.

An iso-WTP curve is illustrated in Figure 1. Point A is the starting drift and volatility, and points B and C are two targets, both of which have the same WTP of $w_1^* = .03$. Moving from A to *any* point on the curve will have the same WTP of .03. Note that α increases as

we move *down* the vertical axis, so the curve represents a social indifference curve between expected return (a lower value of α) and risk (a higher value of σ). It is thus analogous to the return-risk indifference curve of an investor choosing a portfolio of equities and a risk-free asset. The figure also shows another iso-WTP curve, labelled $w_2^* = .05$; moving from α_0 and σ_0 to any point on this curve would have a WTP of .05.

If we knew the cost of reducing α and the cost of reducing σ , we could plot an iso-cost line. If those costs were linear in the change in α and the change in σ , the iso-cost line would be a straight line, as drawn in Figure 1. Its tangency with the iso-WTP curve $w_1^* = .03$ at point B is the cost-minimizing target combination of α and σ . There are other target combinations of α and σ that have a WTP of .03, but they would be more costly to achieve.

In Figure 1, moving from point A to point B has the same WTP = .03 as moving from point A to point C. We can also calculate the combination of *starting values* for α_0 and σ_0 that yield the same WTP = .03 when moving to point B. Doing so is useful because it allows us to determine combinations of temperature drift and volatility under BAU that have the same welfare implications. In other words, how important is the expected change in temperature (under BAU) relative to its variance as determinants of the demand for policy?

This is illustrated in Figure 2, which shows two iso-WTP curves. Curve 1 gives combinations of starting values α_0 and σ_0 such that the movement to point B has a WTP = .03. Thus $w^*(A \rightarrow B) = w^*(D \rightarrow B) = w^*(A \rightarrow C) = .03$. But $w^*(D \rightarrow A) = w^*(B \rightarrow C) = 0$, so $w^*(D \rightarrow C) = .03$. Thus moving from any point on Curve 1 to any point on Curve 2 will have the same WTP.

The iso-WTP and iso-cost curves in Figures 1 and 2 are just illustrative examples. Later I will calibrate this simple model to a set of estimates of temperature change and confidence intervals compiled by the IPCC (2007) and others, and then use the calibrated model to obtain iso-WTP and iso-cost curves along the lines of those shown in Figures 1 and 2.

3 Risk and Return for Policy Outcomes.

Consider a policy to reduce the expected growth rate of temperature and thereby increase the expected growth rate of consumption. Putting aside the cost of the policy, we want to compare the expected benefit to the uncertainty over that benefit, i.e., compare the expected return to the riskiness of that return. This provides a way to compare policies. Two policies might have the same expected benefit, but different degrees of risk.

I begin with a simple two-period example in which consumption today is known but consumption at a future time T is uncertain. I examine the benefit that results from a small increase in the drift of the growth rate (resulting from a small decrease in the drift of temperature). I then turn to a continuous-time model, and again consider a policy that results in a small increase in the drift of the growth rate. Of interest is the expectation and variance of the cumulative return from the policy over the entire horizon.

3.1 Two Periods.

In this simple case, $C_0 = 1$ and C_T is uncertain, the outcome of the stochastic growth process, as given by eqn. (4). Given α and σ , welfare is stochastic and is given by:

$$\tilde{W}_T = \frac{1}{1-\eta} [1 + (C_T^{1-\eta} e^{-\delta T})] \quad (9)$$

Now consider a policy that results (at $t = 0$) in an incremental decrease in α_X , the expected change in temperature, and hence a decrease $\alpha = \gamma\alpha_X$, the expected change in the growth rate of consumption. The (stochastic) return, measured in terms of increased welfare from that policy is

$$\tilde{r}_{\alpha T} = -\frac{\partial \tilde{W}_T}{\partial \alpha} = -\frac{1}{1-\eta} \frac{\partial}{\partial \alpha} C_T^{1-\eta} e^{-\delta T} . \quad (10)$$

Note that

$$C_T^{1-\eta} = e^{(1-\eta) \int_0^T g(s) ds} ,$$

and from eqn. (3), the growth rate $g(s)$ at time $s > 0$ is given by

$$g(s) = g_0 - \alpha s - \sigma \int_0^s dz = g_0 - \alpha s - \sigma z(s) .$$

Thus

$$C_T^{1-\eta} = e^{(1-\eta)g_0T - \frac{1}{2}(1-\eta)\alpha T^2 - \sigma(1-\eta)\int_0^T z(s)ds} , \quad (11)$$

and

$$-\frac{\partial}{\partial\alpha}C_T^{1-\eta} = \frac{1}{2}(1-\eta)T^2C_T^{1-\eta} . \quad (12)$$

The (stochastic) return from this policy is thus

$$\tilde{r}_{\alpha T} = \frac{1}{2}T^2C_T^{1-\eta}e^{-\delta T} . \quad (13)$$

We want the expectation and variance of this return. The Appendix derives a formula for the expectation of $C_T^{1-\eta}$. Using that formula, the expected return is:

$$\bar{r}_{\alpha T} = \frac{1}{2}T^2\mathcal{E}_0(C_T^{1-\eta})e^{-\delta T} = \frac{1}{2}T^2e^{-\rho_0T - \frac{1}{2}\alpha(1-\eta)T^2 + \frac{1}{6}\sigma^2(1-\eta)^2T^3} \quad (14)$$

where

$$\rho_0 \equiv \delta + (\eta - 1)g_0 .$$

Note that $\partial\bar{r}/\partial\alpha > 0$ and $\partial\bar{r}/\partial\sigma^2 > 0$. The first inequality simply results from the fact that the higher is the starting value of α , the greater is the value of reducing α . The second inequality follows from the fact that $\mathcal{E}_0(C_T^{1-\eta})$ is a convex function of the rate of increase in temperature, so that on average stochastic fluctuations in temperature increase the value of reducing its rate of growth. Finally, $\partial\bar{r}/\partial\eta$ can be positive or negative, depending on the relative size of the starting growth rate g_0 .⁷

To get the standard deviation of this return, note from eqn. (12) that

$$\mathcal{E}_0\left(\frac{\partial}{\partial\alpha}C_T^{1-\eta}\right)^2 = \frac{1}{4}(1-\eta)^2T^4\mathcal{E}_0(C_T^{2-2\eta}) . \quad (15)$$

Now use results in the Appendix to get $\mathcal{E}_0(C_T^{2-2\eta})$, so that

$$\mathcal{E}_0\left(\frac{\partial}{\partial\alpha}C_T^{1-\eta}\right)^2 = \frac{1}{4}(1-\eta)^2T^4e^{-2\rho_0T - \alpha(1-\eta)T^2 + \frac{2}{3}\sigma^2(1-\eta)^2T^3} .$$

⁷ $\partial\bar{r}/\partial\eta = [\frac{1}{2}\alpha T^2 - \frac{1}{3}\sigma^2(1-\eta)T^3 - g_0T]\bar{r}_{\alpha T}$. The first two terms in the brackets are positive and the last is negative. If g_0T is large relative to the first two terms, $\partial\bar{r}/\partial\eta$ will be negative. Higher g_0T implies higher C_T and thus greater reduction in marginal utility at T from an increase in η , and hence a lower return from reducing α . A higher value of αT implies the opposite — lower C_T . A higher value of σ^2T^2 implies a higher value of marginal utility at T (via Jensen's inequality).

Thus the variance of the return is given by

$$\begin{aligned}
\mathcal{V}(r_{\alpha T}) &= \mathcal{E}_0(\tilde{r}_{\alpha T}^2) - \bar{r}_{\alpha T}^2 \\
&= \frac{1}{4}T^4 \mathcal{E}_0(C_T^{2-2\eta})e^{-2\delta T} - \bar{r}_{\alpha T}^2 \\
&= \frac{1}{4}T^4 e^{-2\rho_0 T - \alpha(1-\eta)T^2 + \frac{2}{3}\sigma^2(1-\eta)^2 T^3} - \frac{1}{4}T^4 e^{-2\rho_0 T - \alpha(1-\eta)T^2 + \frac{1}{3}\sigma^2(1-\eta)^2 T^3} \\
&= \frac{1}{4}T^4 e^{-2\rho_0 T - \alpha(1-\eta)T^2 + \frac{1}{3}\sigma^2(1-\eta)^2 T^3} [e^{\frac{1}{3}\sigma^2(1-\eta)^2 T^3} - 1] \\
&= \frac{1}{4}T^4 [\mathcal{E}_0(C_T^{1-\eta})e^{-\delta T}]^2 [e^{\frac{1}{3}\sigma^2(1-\eta)^2 T^3} - 1]
\end{aligned} \tag{16}$$

The standard deviation of the return is therefore

$$\text{SD}(r_{\alpha T}) = \frac{1}{2}T^2 \mathcal{E}_0(C_T^{1-\eta})e^{-\delta T} [e^{\frac{1}{3}\sigma^2(1-\eta)^2 T^3} - 1]^{1/2} = \bar{r}_{\alpha T} [e^{\frac{1}{3}\sigma^2(1-\eta)^2 T^3} - 1]^{1/2}. \tag{17}$$

$\text{SD}(r_{\alpha T})$ is proportional to $\bar{r}_{\alpha T}$, so a convenient way to characterize the relative riskiness of this return is by its Sharpe ratio, i.e., the ratio of the expected return to its standard deviation. From eqns. (14) and (17), the Sharpe ratio is

$$S_{\alpha T} = \frac{\bar{r}_{\alpha T}}{\text{SD}(r_{\alpha T})} = [e^{\frac{1}{3}(1-\eta)^2 \sigma^2 T^3} - 1]^{-1/2} \tag{18}$$

The Sharpe ratio is useful because it summarizes the effects of changes in various parameters, such as volatility, σ . As can be seen from eqns. (14) and (17), and increase in σ causes an increase in *both* the expected return and the standard deviation of the return. But the effect on the standard deviation is larger, so that an increase in σ causes a decrease in $S_{\alpha T}$.

Note that the Sharpe ratio is independent of the drift, α . The reason is that the exponential growth rates for both the expected return and the standard deviation include the factor $\frac{1}{2}\alpha(1-\eta)T^2$, which cancels out of the ratio. Also note that if $\eta \neq 1$, $S_{\alpha T} \rightarrow 0$ as $T \rightarrow \infty$. The reason is that the standard deviation of the return grows faster with T than the expected return. Finally, an increase in the index of relative risk aversion, η , leads to a decrease in the Sharpe ratio. The reason is that the return itself is risk-adjusted, i.e., it is proportional not to the change in future consumption but to the change in the *utility* from future consumption. Thus an increase in η has the same effect as an increase in σ .

3.2 Continuous-Time, Finite Horizon.

Now suppose we have a horizon T so that welfare is given by the integral of discounted utility over $t = 0$ to T . Taking future consumption as unknown, welfare is:

$$\tilde{W} = \frac{1}{1-\eta} \int_0^T C_t^{1-\eta} e^{-\delta t} dt \equiv \int_0^T \tilde{W}_t dt . \quad (19)$$

Thus \tilde{W}_t is (uncertain) discounted utility at time $t > 0$.

Once again we consider a policy that increases α incrementally. As before, it is easily seen that the (uncertain) return from this policy is

$$\tilde{r}_\alpha = \int_0^T \frac{1}{2} t^2 C_t^{1-\eta} e^{-\delta t} dt = \int_0^T \frac{1}{2} (1-\eta) t^2 \tilde{W}_t dt . \quad (20)$$

We can now examine returns over the full time horizon.

We want the expectation, standard deviation, and Sharpe ratio for the cumulative return over the horizon 0 to T , \tilde{r}_α . The expectation is straightforward; it is simply

$$\bar{r}_\alpha = \int_0^T \frac{1}{2} t^2 \mathcal{E}_0(C_t^{1-\eta}) e^{-\delta t} dt = \int_0^T \frac{1}{2} t^2 e^{-\rho_0 t - \frac{1}{2} \alpha (1-\eta) t^2 + \frac{1}{6} \sigma^2 (1-\eta)^2 t^3} dt . \quad (21)$$

Note that $\partial \bar{r}_\alpha / \partial \eta$ can be positive or negative depending on the relative magnitudes of g_0 , α and σ , and can change with the time horizon T .

The variance of the return is $\mathcal{V}(\tilde{r}_\alpha) = \mathcal{E}_0(\tilde{r}_\alpha^2) - \bar{r}_\alpha^2$, so we need to find

$$\mathcal{E}_0(r_\alpha^2) = \mathcal{E}_0 \left(\int_0^T \frac{1}{2} t^2 C_t^{1-\eta} e^{-\delta t} dt \right)^2 \quad (22)$$

We can write the integral as:

$$\begin{aligned} r_\alpha^2 &\approx \left[\sum_{j=0}^N \frac{1}{2} j^2 C_j^{1-\eta} e^{-j\delta} \right]^2 \\ &\approx \sum_{i=0}^N \sum_{j=0}^N \frac{1}{4} i^2 j^2 C_i^{1-\eta} C_j^{1-\eta} e^{-\delta(i+j)} \end{aligned} \quad (23)$$

where $N = T/\Delta t$.

As shown in the Appendix,

$$\mathcal{E}_0(C_i^{1-\eta} C_j^{1-\eta}) = C_0^{2(1-\eta)} e^{(1-\eta)g_0(i+j) + b(i,j)} , \quad (24)$$

where

$$b(i, j) = \frac{1}{12}\sigma^2(1 - \eta)^2(i + j)^3 - \frac{1}{2}\alpha(1 - \eta)(i^2 + j^2) . \quad (25)$$

Setting $C_0 = 1$ and letting $\Delta t \rightarrow 0$,

$$\mathcal{E}_0(r_\alpha^2) = \int_0^T \int_0^T \frac{1}{4}t^2s^2 e^{-\rho_0(t+s) - \frac{1}{2}\alpha(1-\eta)(t^2+s^2) + \frac{1}{12}\sigma^2(1-\eta)^2(t+s)^3} dt ds . \quad (26)$$

Using eqns. (21) and (26), we can find the expectation, standard deviation, and Sharpe ratio for the cumulative return \tilde{r}_α . Figure 3 shows the Sharpe ratio as a function of η for two values of σ , a “base case” value of .00012, and .00024, with a time horizon $T = 300$ years. (The values of α and σ are from a calibration to climate change information discussed later.) Figure 4 shows $E(r)$ and $SD(r)$ as functions of η for $\sigma = .00012$. The expected return is decreasing in η ; a higher value of η increases the effective discount rate on consumption. The standard deviation, however, is first increasing and then decreasing in η .

Calculating a Sharpe ratio in this way makes it possible to characterize and compare different policies for different assumptions about parameters such as η and the time horizon T . Two alternative policies might have the same expected return, but the returns might differ considerably in their riskiness.

4 WTP and Policy Targets.

In the previous section I considered marginal policies, i.e., policies that would result in a marginal reduction in α . This allowed for a simple characterization of such policies in terms of a Sharpe ratio. I turn now to policies that would result in a substantial change in the drift α and/or the volatility σ , and evaluate those policies in the context of willingness to pay. I begin with policies for which changes are made at one or more discrete points in time. This is simpler than the continuous-time case, which I turn to next.

4.1 WTP: Consumption Now and at Future T .

Again, $C_0 = 1$, and C_T is uncertain. Under BAU (no abatement policy), $\alpha = \alpha_0$ and $\sigma = \sigma_0$. We want to find the fraction of current and future consumption society would give up to

decrease α and/or decrease σ . Define $a(\alpha_1, \sigma_1, T) = -(1/2)\alpha_1(1-\eta)T^2 + (1/6)\sigma_1^2(1-\eta)^2T^3$. Then to obtain the WTP to move from (α_0, σ_0) to (α_1, σ_1) , write the social welfare for this change as:

$$W_{1T} = \frac{[1 - w(\alpha_1, \sigma_1, T)]^{1-\eta}}{1-\eta} [1 + e^{-\rho_0 T + a(\alpha_1, \sigma_1, T)}] . \quad (27)$$

Under BAU, welfare is

$$W_{2T} = \frac{1}{1-\eta} [1 + e^{-\rho_0 T + a(\alpha_0, \sigma_0, T)}] \quad (28)$$

Equating W_{1T} and W_{2T} yields the WTP for this change:⁸

$$w^*(\alpha_1, \sigma_1, T) = 1 - \left[\frac{1 + e^{-\rho_0 T + a(\alpha_0, \sigma_0, T)}}{1 + e^{-\rho_0 T + a(\alpha_1, \sigma_1, T)}} \right]^{\frac{1}{1-\eta}} , \quad (29)$$

where $\rho_0 \equiv \delta + (\eta - 1)g_0$ as before.

We would like to generate iso-WTP curves, i.e., different combinations of (α', σ') such that $w^*(\alpha', \sigma', T) = w^*(\alpha_1, \sigma_1, T)$. This would describe a “risk-return” trade-off, i.e., a trade-off between decreasing α and decreasing σ that yields the same WTP as for moving from (α_0, σ_0) to (α_1, σ_1) . For this simple two-period case, this is easily done analytically. Let $a(T)$ denote the function a with α and σ set to their starting (calibrated for BAU) values of α_0 and σ_0 . Then using eqn. (29), $w^*(\alpha', \sigma', T) = w^*(\alpha_1, \sigma_1, T)$ implies:

$$a(\alpha', \sigma', T) = a(\alpha_1, \sigma_1, T) , \quad (30)$$

which in turn implies that

$$\sigma' = \sqrt{\sigma_1^2 + \frac{3(\alpha' - \alpha_1)}{(1-\eta)T}} , \quad (31)$$

or equivalently,

$$\alpha' = \alpha_1 + \frac{1}{3}(1-\eta)T(\sigma'^2 - \sigma_1^2) . \quad (32)$$

Thus moving from (α_0, σ_0) to any combination (α', σ') that satisfies eqn. (32) will have the same $WTP = w^*(\alpha_1, \sigma_1, T)$.

⁸Of course it may or may not be possible to shift from (α_0, σ_0) to (α_1, σ_1) . This WTP is simply the most society would sacrifice for this change in α and σ , were it feasible to do so with the resulting revenue.

The slope of the iso-WTP curve is the marginal rate of substitution (MRS) between the two policy targets, α' and σ' . From (32), the MRS is:

$$\text{MRS} = \frac{d\alpha'}{d\sigma'} = \frac{2}{3}\sigma'(1 - \eta)T . \quad (33)$$

Note that as T increases, the MRS becomes larger in magnitude, i.e., it takes larger and larger increases in α' to compensate for any given reduction in σ' . (This is because of the $\sigma^2 T^3$ term versus αT^2 in the function a .)

It is important to keep in mind that the (α', σ') combinations on the iso-WTP curve represent different *ending* positions that result from a policy. In other words, we are always starting from (α_0, σ_0) . All policies that move from this starting position to an ending position on the curve have the same WTP.

4.2 Multiple Time Periods; Action at $t = 0$.

Suppose there are n 50-year time intervals: $T_1 = 50$, $T_2 = 100$, ..., $T_n = 50n$, and suppose we can decrease α and/or reduce σ starting immediately by reducing consumption now (i.e., at $t = 0$) and at every future T_i by some fraction. In this case, to find the WTP to move from (α_0, σ_0) to (α_1, σ_1) , write the welfare measures W_1 and W_2 as:

$$W_1 = \frac{[1 - w(\alpha_1, \sigma_1)]^{1-\eta}}{1 - \eta} \left[1 + \sum_{i=1}^n e^{-\rho_0 T_i + a(\alpha_1, \sigma_1, T_i)} \right] \quad (34)$$

$$W_2 = \frac{1}{1 - \eta} \left[1 + \sum_{i=1}^n e^{-\rho_0 T_i + a(\alpha_0, \sigma_0, T_i)} \right] \quad (35)$$

Thus to obtain $w^*(\alpha_1, \sigma_1)$, just replace eqn. (29) with:

$$w^*(\alpha_1, \sigma_1) = 1 - \left[\frac{1 + \sum_{i=1}^n e^{-\rho_0 T_i + a(\alpha_0, \sigma_0, T_i)}}{1 + \sum_{i=1}^n e^{-\rho_0 T_i + a(\alpha_1, \sigma_1, T_i)}} \right]^{\frac{1}{1-\eta}} . \quad (36)$$

As before, we can generate “iso-WTP” curves, i.e., different combinations of (α', σ') such that $w^*(\alpha', \sigma') = w^*(\alpha_1, \sigma_1)$. These “iso-WTP” curves will depend on n because as we saw above, the $d\sigma^2/d\alpha$ that keeps W_T fixed varies with T . Also, we cannot generate these curves analytically. From eqn. (36), the (α', σ') that gives the same WTP as (α_1, σ_1) is the solution of:

$$\sum_{i=1}^n e^{-\rho_0 T_i + a(\alpha', \sigma', T_i)} = \sum_{i=1}^n e^{-\rho_0 T_i + a(\alpha_1, \sigma_1, T_i)} \quad (37)$$

Eqn. (37) can be solved numerically for combinations of α' and σ' .

Eqn. (37) applies to a specific WTP, namely the WTP that applies to moving from (α_0, σ_0) to (α_1, σ_1) . Starting from (α_0, σ_0) , however, we can find combinations of α' and σ' that yield the same *arbitrary* WTP. Denoting the arbitrary WTP by w^* , we find the combinations of α' and σ' by solving:

$$\sum_{i=1}^n e^{-\rho_0 T_i + a(\alpha', \sigma', T_i)} = (1 - w^*)^{\eta-1} \left[1 + \sum_{i=1}^n e^{-\rho_0 T_i + a(\alpha_0, \sigma_0, T_i)} \right] - 1 . \quad (38)$$

4.3 WTP: Continuous Time with Immediate Action.

Welfare over some future time horizon, with no policy intervention, is given by

$$W_2 = \frac{1}{1 - \eta} \int_0^\infty e^{-\rho_0 t + a(\alpha_0, \sigma_0, t)} dt . \quad (39)$$

Now once again consider a policy that would move from (α_0, σ_0) to (α_1, σ_1) . Denote the WTP for this policy by w_1 , and note that with the intervention, welfare is:

$$W_1 = \frac{1}{1 - \eta} \int_0^\infty (1 - w_1)^{1-\eta} e^{-\rho_0 t + a(\alpha_1, \sigma_1, t)} dt , \quad (40)$$

where once again, $a(\alpha_1, \sigma_1, t) = -(1/2)\alpha_1(1 - \eta)t^2 + (1/6)\sigma_1^2(1 - \eta)^2 t^3$. Equating W_1 and W_2 , we get:

$$w_1 = 1 - \left[\frac{G(\alpha_1, \sigma_1)}{G(\alpha_0, \sigma_0)} \right]^{\frac{1}{\eta-1}} , \quad (41)$$

where

$$G(\alpha_0, \sigma_0) = \int_0^\infty e^{-\rho_0 t + a(\alpha_0, \sigma_0, t)} dt , \quad (42)$$

and likewise for $G(\alpha_1, \sigma_1)$. Thus, given starting values of α and σ (based, say, on a calibration of a future T against the IPCC or other distributions), we can calculate the WTP to increase α and/or decrease σ .

As before, we would like to obtain iso-WTP curves, i.e., combinations of α' and σ' for which the WTP is again w_1 . This can be done numerically by finding combinations of α' and σ' that satisfy

$$G(\alpha', \sigma') = G(\alpha_1, \sigma_1) . \quad (43)$$

We can also obtain combinations of α' and σ' for which the WTP is equal to some arbitrary number, w . From eqn. (41), we find those combinations that satisfy

$$G(\alpha', \sigma') = (1 - w)^{\eta-1} G(\alpha_0, \sigma_0) . \quad (44)$$

Finally, note that the integration over time cannot be carried out over an arbitrarily large horizon. It is necessary to limit the time horizon to some number of years, e.g., 200 or 300.

4.4 Costs of Changing α and σ .

We have seen how we can determine the WTP for moving from a starting drift and volatility, (α_0, σ_0) , to a new (α', σ') , and how we can trace out combinations of α and σ targets that have the same WTP. However, we also want to know the cost of moving from (α_0, σ_0) to (α', σ') . If that cost exceeds the WTP (which is society's reservation price), the policy is economically infeasible. If the cost is less than the WTP, the policy yields a positive surplus.

If we knew the costs of changing α and σ we could plot an iso-cost line. If those costs were linear in the change in α and the change in σ , the iso-cost line would be a straight line, as drawn in Figure 1, and its tangency with an iso-WTP curve would be the cost-minimizing target combination of α and σ . There are other target combinations of α and σ with the same WTP, but they would be more costly to achieve. The minimum cost, however, might be greater or less than the particular WTP. Given a starting point (α_0, σ_0) , the WTP to move to a target (α', σ') depends on the distance to the target and also on the parameters η , δ , and g_0 . For any given shift in α and σ , for example, a reduction in η will increase the WTP, which implies that for any fixed WTP, the required change in α and σ will be smaller, so that the corresponding cost will be smaller.

5 Calibration to Climate Change.

Although it is exceedingly simple, the model specified in Section 2 can be informative about risk and return in the context of climate change policy. In this section, I calibrate the model by determining values of the parameters α_X , σ_X , and γ that are consistent with recent

studies of global warming and its impact, along with estimates of the cost of reducing α_X and σ_X . I then calculate iso-WTP curves and discuss their implications.

5.1 Parameter Values.

To obtain numerical results we need values for α and σ in eqn. (2), the parameters η , δ , and the initial growth rate g_0 , and estimates of the costs of reducing α and σ .

Calibration of α and σ under BAU. The most recent IPCC report (2007) states that growing GHG emissions (i.e., under BAU) would likely lead to a doubling of the atmospheric CO₂e concentration relative to the pre-industrial level by mid-century, which would “most likely” cause an increase in global mean temperature between 2.0°C to 4.5°C by 2100, with an expected value of 2.5°C to 3.0°C. The IPCC report indicates that this range, derived from the results of 22 scientific studies the IPCC surveyed, represents a roughly 66- to 90-percent confidence interval, i.e., there is a 5 to 17-percent probability of a temperature increase above 4.5°C. The 22 studies surveyed by the IPCC also provide rough estimates of the outer tail of the temperature distribution. In summarizing them, the IPCC translated the implied outcome distributions into a standardized form that allows comparability across the studies, and created graphs showing multiple distributions implied by groups of studies. Those distributions suggest that there is a 5% probability that a doubling of the CO₂e concentration relative to the pre-industrial level would lead to a global mean temperature increase, T , of 7°C or more.

I will use, for a horizon of 100 years, $\mathcal{E}(X) = 3^\circ\text{C}$ and a 5% probability of a temperature increase $\geq 7^\circ\text{C}$. The 5% point is 1.65 standard deviations above the mean (because eqn. (1) implies that T_t is normally distributed), so one standard deviation is $4/1.65 = 2.42$. Thus the drift and volatility of the process for temperature are $\alpha_X = 3/100 = .03$, and $\sigma_X = 2.42/\sqrt{100} = 0.242$. More recent studies, e.g., Sokolov et al. (2009), have larger numbers for the expected temperature change and for the 5% point, but I will use these earlier (perhaps overly optimistic) estimates.

We need the process for the growth rate g_t , and thus we need an estimate of γ , which relates g_t to X_t . One way to estimate γ is to use the equation for expected future consumption

derived in the Appendix, combined with estimates of the expected loss of GDP at a specific temperature change. The Appendix shows that expected future consumption is given by:

$$\mathcal{E}_0(C_t) = e^{g_0t - (1/2)\alpha t^2 + (1/6)\sigma^2 t^3} . \quad (45)$$

In earlier work (2011a, 2012), I used the IPCC (2007) estimates that the loss of GDP resulting from a temperature change of 4°C is “most likely” in the range of 1% to 5%. These IPCC estimates are based on several integrated assessment models (IAMs). Taking the upper end of the range, i.e., 5%, implies that with $\alpha'_X = .04$, i.e., $\mathcal{E}_0(X_{100}|\alpha'_X = .04) = 4^\circ\text{C}$, and using eqn. (45), for $t = 100$,

$$\mathcal{E}_0(C_t) = e^{g_0t - (1/2)\alpha'_X \gamma t^2 + (1/6)\sigma_X^2 \gamma^2 t^3} = .95e^{g_0t} ,$$

so that

$$-\frac{1}{2}\alpha'_X \gamma t^2 + \frac{1}{6}\sigma_X^2 \gamma^2 t^3 = \ln(.95) . \quad (46)$$

Substituting $t = 100$, $\alpha'_X = .04$ and $\sigma_X = .242$ yields the following quadratic equation for γ :

$$\gamma^2 - .0205\gamma + .00000526 = 0$$

If $\sigma_X = 0$, $\gamma = .000152$, so we want the smaller root of this equation, which yields $\gamma = .00026$.

This, of course, is an indirect estimate of γ , and relies on the integrated assessment models surveyed by the IPCC that largely posit, rather than estimate, a relationship between temperature and GDP. An alternative way to obtain γ is to utilize direct econometric estimates. Unfortunately, few such estimates are available; the only ones I am aware of are those of Dell, Jones, and Olken (2009, 2012) and Bansal and Ochoa (2011a, b). The latter set of estimates are based on a more extensive data set (147 countries over 1950–2007) and allow for dynamic adjustment of GDP to changes in temperature. Bansal and Ochoa find that a 0.2°C increase in temperature leads to about a .2 percentage point reduction in the growth of real GDP (a number close to that in Dell, Jones and Olken (2009)). However, the Bansal-Ochoa impact number is largely transitory; GDP growth largely recovers after 10 years. In my model, the impact of temperature on GDP growth is permanent, so the Bansal-Ochoa number would imply a value of γ in eqn. (2) of around .001. This is four times

as large as the value obtained from the IPCC’s survey of IAMs, so I will take as “consensus” number a value in the middle, namely $\gamma = .0005$.

Thus we have $\alpha_X = .03$, $\sigma_X = .242$, and $\gamma = .0005$, so that $\alpha = \alpha_X\gamma = .000015$ and $\sigma = \sigma_X\gamma = .00012$. These are the values for α and σ that correspond to the starting values (α_0, σ_0) when calculating WTP and iso-WTP curves.

Other Parameter Values. WTP also depends on the index of relative risk aversion η , the rate of time preference δ , and the base level growth rate g_0 . The finance and macroeconomics literatures have estimates of η ranging from 1.5 to 6, and estimates of δ ranging from .01 to .04. The per capita real growth rate g_0 , measured from historical data, is about .02. It has been argued that for intergenerational comparisons δ should be zero, on the grounds that we should not value the well-being of our great-grandchildren less than our own. Without taking sides in that argument, I set $\delta = 0$ for the calculations that follow. (For values of δ of .01 or higher, WTP is close to zero, and it is difficult to justify *any* policy.)

Likewise, while values of η above 4 may be consistent with the (relatively short-horizon) behavior of investors and consumers, we might apply lower values to welfare comparisons involving future generations. In the (deterministic) Ramsey growth model, the optimal savings rate is $s^* = (R - \delta)/\eta R$, where R is the consumption discount rate and the real return on investment. If $R = .04$ and $\delta = .02$, $s^* = 1/2\eta$, suggesting η should be in the range of 2 to 4. Stern (2007) used a value of 1 for η , which is below the consensus range. Lower values of η result in higher values of WTP; I will generally set $\eta = 2$, but I will also show how various results depend on η .

Costs of Reducing α and σ . Estimating a cost function is speculative at best, given the paucity of available data. Nonetheless, I will proceed on the grounds that some estimate is better than none, and at the very least is illustrative. I will use a linear cost function, i.e., the cost of moving from (α_0, σ_0) to (α', σ') is given by:

$$\text{Cost} = c_1(\alpha_0 - \alpha') + c_2(\sigma_0 - \sigma') \tag{47}$$

Note that the total cost and its two components are expressed in terms of percentages of consumption that must be sacrificed annually. Thus $c_1(\alpha_0 - \alpha')$ is the percentage of

consumption that must be given up each year in order to reduce α from α_0 to $\alpha' < \alpha_0$.

There are no direct estimates of c_1 or c_2 that I am aware of, so coming up with numbers requires extrapolation of the few cost studies that exist, along with some guesswork. I start with the Kyoto Protocol, which if enacted and adhered to is expected to limit any increase in global mean temperature by the end of the century to roughly 3°C (which is the *expected* temperature increase under BAU). The U.S. Energy Information Administration (1998) estimated that compliance with the Protocol would cost from 1 to 3% of GDP annually. Estimates from country cost studies in IPCC (2007c) also put the cost of limiting the increase in temperature to 3°C at around 1 to 3% of GDP. I will take the mid-range of 2% of GDP and translate “limiting the increase in temperature to 3°C ” to mean reducing the expected temperature increase from 3°C to 1.5°C with a standard error of about 1.5°C . This means cutting the drift α in half, i.e., from .000015 to .0000075, which in turn implies that $c_1 = (.02)/(7.5 \times 10^{-6}) = 2667$.

Obtaining an estimate for c_2 , the cost of reducing the volatility and hence the standard deviation of future temperature change, is even more difficult. According to numbers compiled by Palmer (2011), public funding for climate change research in the U.S. during 2009 was about \$5 billion. These expenditures could arguably improve our understanding of global warming and reduce our uncertainty over future temperatures; they exclude public R&D spending on renewable energy, new technologies, electricity and fossil fuel production, and related technologies.⁹ Also, this ignores R&D spending directed at adaptation to climate change (which so far has been very limited). I will make the bold assumption that a doubling of the \$5 billion in expenditures to \$10 billion per year could result in a halving of the volatility σ_0 . In other words, reducing σ from .00012 to .00006 would require a sacrifice of about 0.067% of GDP annually (based on a \$15 trillion GDP), which in turn implies that $c_2 = .00067/.00006 = 11.17$. This number should be viewed as a “guesstimate,” but it is illustrative, and I will also examine the implications of alternative values for c_2 .

⁹For the six focus areas of The US Global Change Research Program, this figure includes only “Improving Knowledge,” “Improving Understanding,” and “Modelling and Prediction.” It also includes about \$1 billion of science expenditures under the 2009 Recovery Act that are climate related. The U.S. GAO (2003) puts the 2003 expenditure at about \$4 billion.

5.2 Numerical Results.

I begin by calculating iso-WTP curves using the parameter values discussed above. Figure 5 shows three iso-WTP curves, generated using the starting values $\alpha_0 = .000015$ and $\sigma_0 = .00012$ based on the calibration described above. Other parameters are $g_0 = .02$, $\delta = 0$, $\eta = 2$, and the time horizon is $T = 300$ years. The starting values (α_0, σ_0) are shown as the small circle near the bottom of the diagram. The solid line in the middle is the iso-WTP curve for $\text{WTP} = .0365$, which is the WTP to reduce α to zero but leave σ unchanged. Thus moving from (α_0, σ_0) to any point on that curve has a WTP of .0365. The other two curves in Figure 5 show combinations of α and σ that, starting from α_0 and σ_0 , have WTPs of .02 and .06. (Note that all of the values of α on the top curve are negative.)

Figure 6 shows an iso-WTP curve (labeled B) for the same starting and ending values of the drift and volatility, $(\alpha_0, \sigma_0) = (.000015, .00012)$ and $(\alpha', \sigma') = (0, .00012)$, and the same values for the other parameters, so that the WTP is again .0365. The curve labeled A shows alternative combinations of starting values for α_0 and σ_0 that give the same $\text{WTP} = .0365$ for any target on curve B. Thus shifts from any points on curve A to any points on curve B have a WTP of .0365 and thus are welfare-equivalent.

The iso-WTP curves in Figure 5 show the most that society would be willing to pay (in terms of percentage of consumption sacrificed) in order to move from a starting point (α_0, σ_0) to an ending point on the curve. They do not, however, tell us anything about the cost of moving from (α_0, σ_0) to an ending point. If the cost exceeds the WTP, the policy is economically infeasible; if the cost is less than the WTP, the policy has a positive social surplus. In addition, since all points on an iso-WTP curve are welfare-equivalent, we might like to know which point has the least cost. To address these questions I utilize the iso-cost lines implied by eqn. (47), with the calibrated values for c_1 and c_2 , i.e., the cost equation

$$\text{Cost} = 2667(\alpha_0 - \alpha') + 11.17(\sigma_0 - \sigma') . \quad (48)$$

Figure 7 shows an iso-WTP curve that plots combinations of target values of α and σ that have a WTP of .02, given the starting point $(\alpha_0, \sigma_0) = (.000015, .00012)$, which is shown as a small circle near the bottom of the figure. Also shown is an iso-cost line that is tangent

to the iso-WTP curve, i.e., at the point where the slope of the iso-WTP curve is $-c_2/c_1$. (Remember that α is *decreasing* as we move up the vertical axis.) Thus at this point the cost of moving from (α_0, σ_0) to the iso-WTP curve is minimized. In this example, that cost turns out to be about .020, i.e., the same as the WTP. Thus the policy illustrated by Figure 7 is just feasible, i.e., society would be paying its reservation price for the policy, so the social surplus would be zero. On the other hand, if we were to double both c_1 and c_2 (so that the slope and tangency point remained the same), the cost would also double, making the policy economically infeasible. And if we were to instead halve both c_1 and c_2 , the cost would drop to .010, so that the policy would have a positive social surplus of .01, i.e., 1% of GDP.

Given eqn. (48), the cost of any policy is a simple function of the changes in α and σ . But the WTP for changes in α and σ depends not only on the magnitudes of the changes, but also on the parameters η , δ , and g_0 . For any given shift in α and σ , for example, a reduction in η will increase the WTP.¹⁰ Equivalently, a lower value of η implies that for any fixed WTP, the required changes in α and σ will be smaller. This is illustrated by Figure 8, which is the same as Figure 7 except that $\eta = 1.5$ instead of 2. If $\eta = 1.5$ it takes much smaller reductions in α and σ to have a WTP of .02, and thus the cost of the policy is much smaller (.008 versus .020). In this case the policy has a positive social surplus.

Table 1 shows the target (α', σ') , the cost, and the net social surplus for other parameter values. The first row shows the base case corresponding to Figure 7, and the second row corresponds to Figure 8 (in which η is reduced to 1.5 but the WTP is fixed at .02, so that the cost falls and the social surplus is positive). Increasing η to 3 implies that for any fixed target (α', σ') , the WTP will fall, so that obtaining a WTP of .02 requires a much larger shift in (α', σ') , and thus a much higher cost (in this case .091 so that the social surplus is negative). As the fourth row shows, with $\eta = 3$ we can still obtain a positive social surplus, but only when the WTP is very small. As the fifth row of the table shows, reducing g_0 has the same effect as reducing η (the initial deterministic consumption discount rate is $R_0 = \delta + \eta g_0$), so

¹⁰Reducing η reduces the effective consumption discount rate, which, ignoring uncertainty, is given by $R_t = \delta + \eta g_t$ (the deterministic Ramsey rule). Thus reducing η increases the present value of future losses of consumption resulting from higher temperatures. For further discussion of this point, see Pindyck (2012).

Table 1: WTP AND COST

Note: Entries show changes from the base case, for which $c_1 = 2667$, $c_2 = 11.17$, $\eta = 2$, $g_0 = .02$, and $\delta = 0$. Throughout, starting point is $\alpha_0 = .000015$ and $\sigma_0 = .00012$.					
Change from Base Case	WTP	α'	σ'	Cost	Social Surplus
BASE CASE	.02	.0000077	.000046	.020	0
$\eta = 1.5$.02	.000012	.000068	.008	.012
$\eta = 3$.02	-.0000187	.000046	.091	-.071
$\eta = 3$.00032	.0000152	.000041	.00032	0
$g_0 = .01$.02	.0000131	.000033	.0060	.014
$\delta = .01$.02	-.0000031	.000067	.049	-.029
$c_2 = 22.34$.02	.0000074	.000090	.021	-.001

that the target (α', σ') is closer to the initial (α_0, σ_0) and the cost falls. Increasing δ has the same directional effect as increasing η , so that with WTP fixed at .02, the cost increases. Finally, in the last row of the table, c_2 , the cost of reducing σ , is doubled. As we would expect, this increases the overall cost (but not by much), and causes the target (α', σ') to shift to a larger reduction in α and smaller reduction in σ .

5.3 Policy Design and Evaluation.

Taking this simple model and the calibrated parameter values at face value, we can infer the optimal mix of risk and return as targets of climate change policy. We can likewise determine whether a particular pair of policy targets (which might or might not be an optimal mix) yields a positive social surplus, i.e., has a cost less than the WTP for the targets. Finally, we can determine alternative target combinations that are welfare-equivalent, in that they lie on the same iso-WTP curve.

In Figure 7, for example, the iso-WTP curve corresponds to a WTP of .02, and given the starting point (i.e., α_0 and σ_0 under BAU), the optimal target mix calls for cutting both α and σ roughly in half. The cost of reaching this optimal target mix turns out to be .02 as well, so that the policy is just feasible. We could have instead chosen an alternative target mix that also has a WTP of .02 (i.e., a different point on the iso-WTP curve), but the cost

of that policy would exceed .02.

As is evident from Figure 8 and Table 1 the social surplus for any particular policy depends not only on the iso-WTP curve and the cost parameters c_1 and c_2 , but also on the parameters δ , η , and g_0 . The dependence on δ and η in particular has been a long-standing problem for the design and evaluation of climate policy. What are the “correct” values for these parameters that should be used for policy evaluation? Economists disagree. The answer depends in part on whether we view these as behavioral parameters (i.e., reflecting the behavior of consumers, investors, and firms) or policy parameters (i.e., reflecting the opinions and objectives of policy makers). Either way, we are left with a wide range of reasonable values, and thus a wide range of estimates of WTP for any target mix. As a comparison of Figures 8 and 8 show, a value of η of 1.5 versus 2.0 yields very different WTPs for any given set of targets (or, as in the figures, very different targets for a fixed WTP). The approach to policy design outlined in this paper does not get around this problem. Thus it is probably best used as a way of clarifying policy target tradeoffs (risk versus return) conditional on some set of values for η , etc.

6 Concluding Remarks.

Using a simple model of a stock externality (e.g., temperature) that evolves stochastically, I examined the extent to which the value of a costly policy (e.g., emissions abatement) depends on the policy’s expected future benefits versus their riskiness. I also presented a framework by which one can estimate the extent to which the objective of policy should be a reduction in the expected future damages versus a reduction in the uncertainty over those damages. The framework is based on the computation of “iso-WTP curves” (social indifference curves) for combinations of risk and expected returns as policy objectives. Given cost estimates for reducing risk and increasing expected returns, one can then compute the optimal risk-return mix for policy, and evaluate the policy’s social surplus.

To illustrate this framework, I used a simple Brownian motion process for the stock variable, in this case temperature, and a CRRA utility function. This has the advantage

of yielding simple expressions for expected future consumption and utility. However, the framework can be applied to alternative stochastic processes for the stock variable and alternative social preferences, although numerical methods may be needed to obtain results.

In the case of climate change, there is a broad consensus on the expected rate of temperature increase under BAU — about 3°C. There is considerable disagreement, however, over the nature and extent of uncertainty around that expected value. That disagreement relates to the choice of probability distribution that best applies to future temperatures (e.g., whether the distribution should be fat- or thin-tailed), and to the parameterization of any particular distribution.¹¹ I have used a simple (thin-tailed) normal distribution to describe temperature change, but one can incorporate alternative distributions (using alternative stochastic processes for X_t). Thus the framework presented here can be used to evaluate the implications of alternative characterizations of uncertainty for determining the targets of climate change policy.

It is an understatement to say that caveats are needed. First, costs of reducing α and σ are probably convex, not linear as in eqn. (47). But a linear cost function may be a reasonable first approximation, and in any case our ability to estimate these costs is limited. More troublesome is the assumption that whatever the cost, we can change α and σ instantly. One would expect that shifts in α and σ would occur gradually in response to ongoing expenditures on abatement and R&D. Also, I calculated WTP under the restriction that the flows of costs are constant percentages of GDP. In fact, those costs might fall as a result of technological change making it cheaper to limit GHG emissions. And even if the costs do not fall, WTP might be calculated under the alternative assumption that costs increase or decrease over time as the policy becomes more or less stringent. The model can be modified to allow for gradual changes in α and σ and/or costs that are time-varying (although both changes would be at the cost of complicating the numerical solution of the model).

Even maintaining the simple assumptions used here (in which the change in α and σ occur instantly and the flow cost is a constant percentage of GDP), the framework can be

¹¹Pindyck (2011b) addresses the question of fat- versus thin-tailed distributions in the context of climate change policy.

extended in a number of directions. For example, the timing of policy implementation can be made endogenous. Suppose a policy is imposed only when temperature reaches a critical threshold (e.g., 3°C). Given a cost for the policy, we could then determine the temperature threshold that maximizes net welfare.

Appendix: Some Expectations

There are some key expected values that are used throughout this paper. They are derived below, beginning with $\mathcal{E}_s(C_t^{1-\eta})$.

Expected Utility. Denote $F(C, g, s) = \mathcal{E}_s(C_t^{1-\eta})$, for $t > s$. Then F must satisfy the following partial differential equation (i.e., Kolmogorov Forward Equation):

$$dF = \frac{1}{2}\sigma^2 F_{gg} - \alpha F_g + gCF_C + F_s = 0 , \quad (49)$$

with boundary conditions $F(C, g, t) = C_t^{1-\eta}$ and $F(0, g, s) = 0$. Try, and then verify by substitution, a function of the form:

$$F(C, g, s) = C_s^{1-\eta} e^{(1-\eta)g_s(t-s)+a(s)} . \quad (50)$$

The partial derivatives are:

$$\begin{aligned} F_g &= (1-\eta)(t-s)F \\ F_{gg} &= (1-\eta)^2(t-s)^2F \\ F_C &= (1-\eta)F/C \\ F_s &= [-(1-\eta)g_s + a'(s)]F \end{aligned}$$

After substituting into eqn. (49) and dividing through by F , we have:

$$a'(s) = -\frac{1}{2}\sigma^2(1-\eta)^2(t-s)^2 + \alpha(1-\eta)(t-s) \quad (51)$$

so that the function $a(s)$ is:

$$a(s) = \frac{1}{6}\sigma^2(1-\eta)^2(t-s)^3 - \frac{1}{2}\alpha(1-\eta)(t-s)^2 . \quad (52)$$

Note that if $\eta > 1$ and $\alpha > 0$, $a(s) > 0$.

Thus starting at time $s = 0$ with $C_0 = 1$,

$$\mathcal{E}_0(C_t^{1-\eta})e^{-\delta t} = e^{-\delta t + (1-\eta)g_0 t - (1/2)\alpha(1-\eta)t^2 + (1/6)\sigma^2(1-\eta)^2 t^3} . \quad (53)$$

Note that if $\eta > 1$, as t increases, $\mathcal{E}_0(C_t^{1-\eta})$ first decreases and then eventually increases without bound. Thus the integration over time to compute welfare must be done for a finite horizon.

Expected Consumption. We will also need an expression for $\mathcal{E}_0(C_t)$. Following the same steps as above it is easy to show that if $C_0 = 1$,

$$\mathcal{E}_0(C_t) = e^{g_0 t - (1/2)\alpha t^2 + (1/6)\sigma^2 t^3} . \quad (54)$$

Note that if $\alpha^2 > 2g_0\sigma^2$, $\mathcal{E}_0(C_t)$ is first increasing in t up to a local maximum, then will decrease to a local minimum, and then increase without bound. But if $\alpha^2 < 2g_0\sigma^2$, then $\mathcal{E}_0(C_t)$ will always be increasing in t .

Expectations of Products. We will want to calculate the variance of the return from an abatement policy. To do this we need the expectation of $C_i^{1-\eta}C_j^{1-\eta}$ for two different future times i and j . Denote $G(C, g, s) = \mathcal{E}_s(C_i^{1-\eta}C_j^{1-\eta})$ for $i, j > s$. Then G satisfies the following Kolmogorov Forward Equation:

$$dG = \frac{1}{2}\sigma^2 G_{gg} + \alpha G_g + gCG_C + G_s = 0 , \quad (55)$$

Try and confirm that the solution has the form:

$$G(C, g, s) = C_s^{2(1-\eta)} e^{(1-\eta)g_s[(i-s)+(j-s)]+b(s)} . \quad (56)$$

Substituting the partial derivatives G_g, G_{gg} , etc. into eqn. (55) yields

$$b'(s) = -\frac{1}{2}\sigma^2(1-\eta)^2[(i-s)+(j-s)]^2 + \alpha(1-\eta)[(i-s)+(j-s)] \quad (57)$$

so that the function $b(s)$ is:

$$b(s) = \frac{1}{12}\sigma^2(1-\eta)^2[(i-s)+(j-s)]^3 - \frac{1}{2}\alpha(1-\eta)[(i-s)^2+(j-s)^2] . \quad (58)$$

I use this result when calculating the variance of incremental returns from an abatement policy.

References

- Bansal, Ravi, and Marcelo Ochoa, “Welfare Costs of Long-Run Temperature Shifts,” NBER Working Paper 17574, Nov. 2011a.
- Bansal, Ravi, and Marcelo Ochoa, “Temperature, Aggregate Risk, and Expected Returns,” NBER Working Paper 17575, Nov. 2011b.
- Barrett, Scott, “The Incredible Economics of Geoengineering,” unpublished manuscript, March 2007.
- Brock, William A., and M. Scott Taylor, “The Green Solow Model,” *Journal of Economic Growth*, 2010.
- Dasgupta, Partha, “Discounting Climate Change,” *Journal of Risk and Uncertainty*, 2008, **37**, 141–169.
- Debreu, Gerard, “A Classical Tax-Subsidy Problem,” *Econometrica*, Jan. 1954, **22**, 14–22.
- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken, “Temperature and Income: Reconciling New Cross-Sectional and Panel Estimates,” *American Economic Review*, May 2009, **99**, 198–204.
- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken, “Temperature Shocks and Economic Growth: Evidence from the Last Half Century,” *American Economic Journal: Macroeconomics*, 2012, **4**, 66–95.
- Energy Information Administration, “Comparing Cost Estimates for the Kyoto Protocol,” Report No. SR/OIAF/98-03, Oct. 1998.
- Heal, Geoffrey, “Climate Economics: A Meta-Review and Some Suggestions,” *Review of Environmental Economics and Policy*, Winter 2009, **3**, 4–21.
- Heal, Geoffrey, and Bengt Kriström, “Uncertainty and Climate Change,” *Environmental and Resource Economics*, June 2002, **22**, 3–39.
- Intergovernmental Panel on Climate Change (IPCC), “Climate Change 2007: The Physical Science Basis,” Cambridge University Press, 2007a; “Climate Change 2007: Impacts, Adaptation, and Vulnerability,” 2007b; “Climate Change 2007: Mitigation of Climate Change,” 2007c.
- Nordhaus, William D., *A Question of Balance: Weighing the Options on Global Warming Policies*, Yale University Press, 2008.
- Palmer, Lisa, “Show Me the Money,” *Nature Climate Change*, October 2011, **1**, 376–380.
- Pindyck, Robert S., “Uncertainty in Environmental Economics,” *Review of Environmental Economics and Policy*, Winter 2007.

- Pindyck, Robert S., “Modeling the Impact of Warming in Climate Change Economics,” in G. Libecap and R. Steckel, eds., *The Economics of Climate Change: Adaptations Past and Present*, University of Chicago Press, 2011a.
- Pindyck, Robert S., “Fat Tails, Thin Tails, and Climate Change Policy,” *Review of Environmental Economics and Policy*, Summer 2011b, **5**, 258–274.
- Pindyck, Robert S., “Uncertain Outcomes and Climate Change Policy,” *Journal of Environmental Economics and Management*, February 2012.
- Roe, Gerard H., and Marcia B. Baker, “Why is Climate Sensitivity So Unpredictable?” *Science*, October 26, 2007, **318**, 629–632.
- Sokolov, A.P., P.H. Stone, C.E. Forest, R. Prinn, M.C. Sarofim, M. Webster, S. Paltsev, C.A. Schlosser, D. Kicklighter, S. Dutkiewicz, J. Reilly, C. Wang, B. Felzer, and H.D. Jacoby, “Probabilistic Forecast for 21st Century Climate Based on Uncertainties in Emissions (without Policy) and Climate Parameters,” MIT Joint Program on the Science and Policy of Global Change, Report No. 169, January 2009.
- Stern, Nicholas, *The Economics of Climate Change: The Stern Review*, Cambridge University Press, 2007.
- Stokey, Nancy, “Are There Limits to Growth?” *International Economic Review*, Feb. 1998, **39**, 1–31.
- U.S. Government Accountability Office, “Federal Climate Change Expenditures: Report to Congress,” GAO-05-461, August 2003.
- Weitzman, Martin L., “On Modeling and Interpreting the Economics of Catastrophic Climate Change,” *Review of Economics and Statistics*, February 2009a.
- Weitzman, Martin L., “Additive Damages, Fat-Tailed Climate Dynamics, and Uncertain Discounting,” *Economics: The Open-Access E-Journal*, 2009b, **3**, 39.




Figure 1: Iso-WTP Curves and Iso-Cost Line




Figure 2: Iso-WTP Starting Curve and Target Curve




Figure 3: Sharpe Ratio vs. η ($g_0 = .020$, $\delta = 0$, $T = 300$, $\sigma = .00012$ and $.00024$)




Figure 4: Mean and Standard Deviation of Return vs. η ($g_0 = .020$, $\delta = 0$, $T = 300$, $\sigma = .00012$)




Figure 5: Iso-WTP Curves ($g_0 = .020$, $\delta = 0$, $\eta = 2$, $T = 300$)




Figure 6: Iso-WTP Starting and Target Curves ($g_0 = .020$, $\delta = 0$, $\eta = 2$, $T = 300$)




Figure 7: Iso-WTP and Iso-Cost Curves ($g_0 = .020$, $\delta = 0$, $\eta = 2$, $T = 300$)

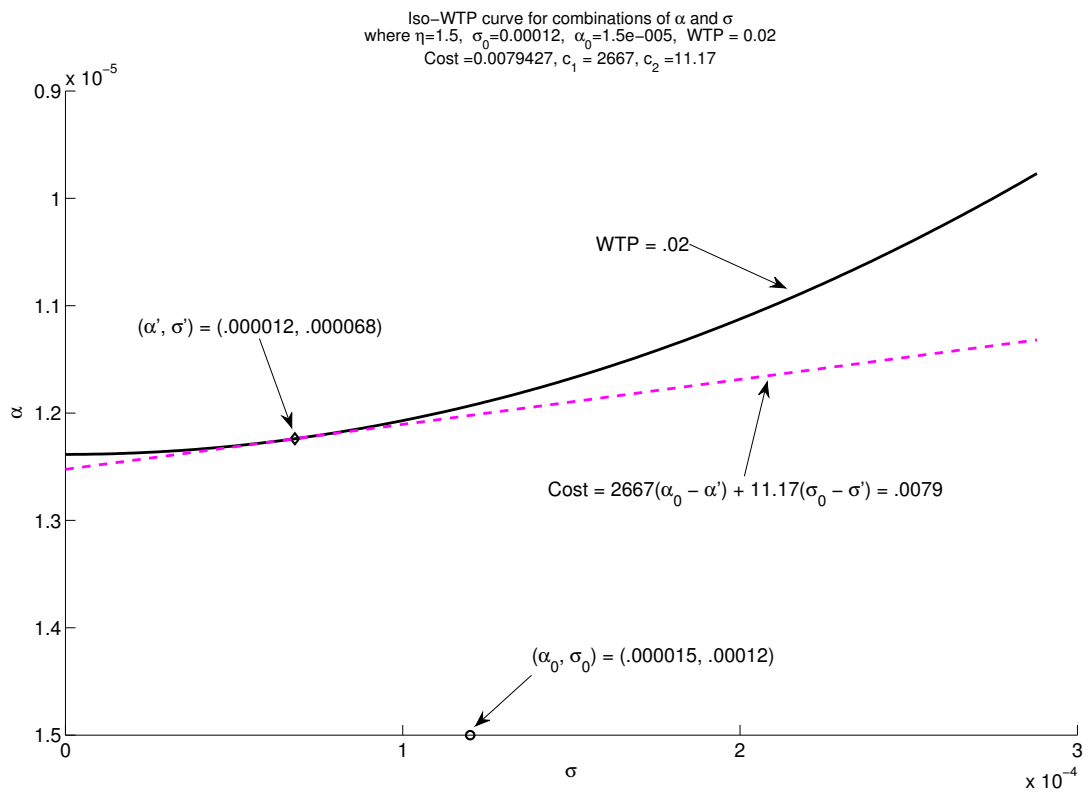


Figure 8: Iso-WTP and Iso-Cost Curves ($g_0 = .020$, $\delta = 0$, $\eta = 1.5$, $T = 300$)