Projects to Avoid Catastrophes

Robert S. Pindyck*

April 10, 2013

Abstract

How should we evaluate public policies or projects that are intended to reduce the likelihood or potential impact of a catastrophic event? Examples might include a greenhouse gas abatement policy to avert a climate change catastrophe, or the construction of levees to avert major flooding. Projects that reduce the likelihood or potential impact of catastrophic events are essentially out-of-the-money put options, can be evaluated as such, and can have negative discount rates. However, the value of the project (as a contingent claim) may be higher or lower than society's willingness to pay for it, and may be higher or lower than the government's ability to actually provide it.

1. Introduction.

Suppose that over the next 50 years there is some likelihood that climate change will have a catastrophic impact on GDP and consumption. Suppose that by allocating some fraction of current and future GDP to emissions abatement, we could reduce or eliminate the chance of this catastrophic outcome, or reduce its likely impact should it occur. How should we evaluate this project and what discount rate should we use? Can we view the project as essentially an out-of-the-money put option, so that the appropriate discount rate may be negative?

We can think of the project as follows. There is some possibility that under "business as usual" (BAU, i.e., no project), there is some probability that a catastrophic event will occur that would reduce GDP and consumption by some substantial amount. If, on the other hand, society gives up a fraction of current and future consumption, the government will in return guarantee a much lower expected cost should the catastrophic

*Sloan School of Management, M.I.T., Cambridge, MA. My thanks to Deborah Lucas for helpful comments and suggestions.

event occur (or guarantee that the catastrophic event will not occur). At this point, *how* the government manages to provide this guarantee is not important. In the case of climate change, one might imagine that the money the government collects will be used to reduce GHG emissions, or finance some technology ("geoengineering") to reverse the effects of rising atmospheric GHG concentrations. We will simply assume that the government is able to provide the guarantee, given the money it collects.

Can we therefore value the project in the same way that we would value a loan guarantee? Of course unlike a loan guarantee (where the guarantor spends money only if the bad outcome occurs), in this case the government might have to spend the savings from reduced consumption immediately, rather than wait to see whether the catastrophic outcome indeed occurs. But that immediate expenditure (should it be needed) occurs "behind the scenes." In other words, from the public's point of view, a payment is being made now to avoid a bad future outcome, and exactly what the government does to avoid the bad outcome doesn't really matter.

This note begins with a simple two-period example, based largely on Lucas (2002, 2013), to explore this approach to project evaluation. The project is a guarantee to society that there will be no catastrophic outcome. Without the project (i.e., under BAU) there is a known probability that the catastrophic outcome will occur, and will reduce future consumption by a known fraction. I calculate the value of this guarantee in terms of a fraction of current and future consumption.

As with most call or put options, the value of the guarantee does not depend on the underlying probability and impact of a catastrophe under BAU. The value only depends on the normal growth rate of consumption, the risk-free interest rate, and the length of time for which the guarantee applies. With a typical financial guarantee we can infer the probability of a bad outcome from the expected return on the underlying asset, consistent with the CAPM or other asset pricing model. In this case, the probability and impact will have to be inferred from a general equilibrium model of consumption growth, interest rates, etc. Furthermore, the probability and impact may be such that society would not want to "buy" the guarantee, i.e., it may be that total expected utility is higher without the guarantee.

The two-period aspect of this example is quite restrictive. I therefore lay out a continuous-time model in which the value of the project is determined in a general equilibrium context. This model is essentially a stripped-down version of the model presented in Pindyck and Wang (forthcoming).

2. A Simple Two-Period Example.

We will be concerned with consumption at only two dates — today ($t = 0$) and a future date T, which I will initially take to be 50 years from now. Consumption today is C_0, and barring a catastrophe will grow at the "normal" annual rate g, so that over T years it will have grown by $(1+g)^T - 1 \equiv G_T$. Under BAU, there is a probability Λ that a catastrophe will occur on or before time T; if it does occur, consumption at time T will be reduced by a fraction ϕ. Thus under BAU consumption evolves as:

$$
C_0 \quad \begin{array}{lll} \nearrow & (1+G_T)C_0 & \text{probability } 1-\Lambda \\ \searrow & (1-\phi)(1+G_T)C_0 & \text{probability } \Lambda \end{array}
$$

Thus expected future consumption is $\mathcal{E}_0(C_T) = (1-\phi\Lambda)(1+G_T)C_0$. This implies that the T-period expected rate of growth of consumption (allowing for a possible catastrophe) is $G_T' = (1-\phi\Lambda)(1+G_T) - 1$, which is an annual rate of $g_c' = [(1-\phi\Lambda)(1+G_T)]^{1/T} - 1$.

Suppose a guarantor (the government) is paid a fraction θ of current and future consumption, and in return will eliminate the risk of the bad outcome.[1] Then current consumption is only $(1-\theta)C_0$, and will grow to $C_T = (1+G_T)(1-\theta)C_0$. What is the value of this guarantee? Again, *how* the guarantor manages to eliminate the risk is not of concern; we assume that doing so is feasible, and just want to know how valuable the guarantee is. Then we can think of the guarantor as agreeing to replace the loss of consumption under the bad outcome. That loss of consumption is $\phi(1+G_T)(1-\theta)C_0$. Thus the guarantor's payoffs under the two possible outcomes are:

$$
P \quad \begin{array}{lll} \nearrow & 0 & \text{probability } 1-\Lambda \\ \searrow & -\phi(1+G_T)(1-\theta)C_0 & \text{probability } \Lambda \end{array}
$$

The expected loss to the guarantor is $\phi\Lambda(1+G_T)(1-\theta)C_0$.

Finally, we assume there is a risk-free T-year bond, with an annual risk free rate of r_f, and I will assume that $r_f < g$. Then the risk-free rate over T years is $(1+r_f)^T - 1 \equiv R_f$. The current and future values of the bond are given by:

$$
B_0 = \frac{100}{1+R_f} \quad \begin{array}{ll} \nearrow & 100 \\ \searrow & 100 \end{array}
$$

[1]The claim on consumption can can occur through a simple consumption tax, or can be a claim on the capital stock. For example, with an AK production technology, the government would own a fraction of K. See, e.g., Pindyck and Wang (forthcoming).

As in Lucas (2002), create a portfolio that replicates the possible payoffs from the guarantee, so that its value must equal the value of the guarantee. The portfolio has a claim on n units of consumption and is short m units of the risk-free bond. Thus its current value is:

$$\Phi_0 = nC_0 - mB_0 = nC_0 - \frac{100m}{1+R_f} \tag{1}$$

The value of the portfolio at time T is:

$$\Phi_T = \begin{cases} (1+G_T)nC_0 - 100m & \text{if}\, C_T = (1+G_T)C_0 \\ \\ (1-\phi)(1+G_T)nC_0 - 100m & \text{if}\, C_T = (1-\phi)(1+G_T)C_0 \end{cases}$$

Choose n and m so that Φ_T replicates the guarantee. Thus:

$$(1+G_T)nC_0 - 100m = 0 \tag{2}$$

$$(1+G_T)(1-\phi)nC_0 - 100m = -(1-\theta)(1+G_T)\phi C_0 \tag{3}$$

Solving these two equations for n and m yields

$$n = 1-\theta \qquad \text{and} \qquad m = \frac{(1+G_T)(1-\theta)}{100}C_0$$

Plugging these values for n and m into eqn. (1) gives:

$$\Phi_0 = (1-\theta)C_0\left[1 - \frac{1+G_T}{1+R_f}\right] \;. \tag{4}$$

By assumption, $R_f < G_T$, so $\Phi_0 < 0$. Thus the value of the guarantee (i.e., what the guarantor must be paid) is

$$-\Phi_0 = \left[\frac{G_T - R_f}{1+R_f}\right](1-\theta)C_0 \;. \tag{5}$$

To find θ, set this payment equal to θC_0:

$$\theta = \frac{G_T - R_f}{1+G_T} \;. \tag{6}$$

Thus the value of the guarantee is:

$$-\Phi_0 = \frac{G_T - R_f}{1+G_T}C_0 = \theta C_0 \,. \tag{7}$$

Note that the value of the guarantee (and the fraction of consumption θ allocated to the guarantee) are independent of the probability of a catastrophe, Λ, and its impact ϕ. It might be that the values of Λ and ϕ can be inferred from the behavior of economic variables (as in Pindyck and Wang (forthcoming)), but I will simply assume that these values are exogenous and given.

2.1. The Effective Discount Rate.

What is the effective discount rate for this "project?" Recall that the expected payment by the guarantor is $\phi\Lambda(1+G_T)(1-\theta)C_0$. Using eqn. (6) for θ, the expected payment is $\phi\Lambda(1+R_f)C_0$. Letting R_T denote the effective discount rate over T years and using eqn. (7) for the value of the guarantee:

$$\frac{\phi\Lambda(1+R_f)C_0}{1+R_T} = -\Phi_0 = \left[\frac{G_T - R_f}{1+G_T}\right] C_0 \tag{8}$$

Thus the discount rate is:

$$R_T = \frac{\phi\Lambda(1+R_f)(1+G_T)}{G_T - R_f} - 1\ . \tag{9}$$

2.2. Some Numbers.

Let's introduce some plausible numbers. Again, we will use a 50-year time horizon. The risk-free rate in the U.S. over the past 60 years has been roughly $r_f = .01$, so $R_f = 1.01^{50} - 1 = 0.64$. The average real per-capital growth rate has been roughly $g = .02$, so $G_T = 1.02^{50} - 1 = 1.69$. Lastly, we set $\Lambda = \phi = 0.3$.

These numbers imply that $\theta = 0.39$ and the value of the guarantee is $0.39C_0$. (Again, this is independent of Λ and ϕ.) What is the effective discount rate for the guarantee? Using eqn. (9), over the 50-year horizon it is $R_T = -0.622$. This implies an annual discount rate of -0.0193, i.e., close to negative 2%. The negative discount rate is warranted because this "project" shifts risk over to the guarantor. Put another way, it is the effective rate of return on a put option (which society exercises only if the bad outcome occurs). That rate of return is typically negative.

Why is the value of the guarantee so large — 39 percent of consumption? The reason is that it is guaranteeing a fixed rate of consumption growth (2%) over a very long time horizon, 50 years. Suppose we repeat this exercise for a 10-year horizon. Assuming $r_f = .01$ and $g = .02$ as before, the 10-year rates are now $R_f = 1.01^{10} - 1 = 0.10$ and $G_T = 1.02^{10} - 1 = 0.22$. From eqn. (6), we now have $\theta = 0.10$, i.e., the guarantee is now worth only 10 percent of consumption. Once again, the value of the guarantee is independent of the probability or impact of a catastrophic event.

The value of the guarantee depends on g and r_f. Suppose $g = .015$ instead of .02, so the 50-year growth rate is $G_T = 1.015^{50} - 1 = 1.105$. Then $\theta = 0.22$ and the value of the guarantee is $0.22C_0$.

2.3. Willingness to Pay.

Would society sacrifice a substantial fraction of consumption for a guarantee that maintains a normal growth rate? To address this, we need to specify a social utility function. I will use the CRRA utility

$$u(C) = \frac{1}{1-\eta} C^{1-\eta} \ ,$$

and I assume $\eta > 1$. Let δ be the annual rate of time preference. To compare guarantees for different time horizons, we must account for the fact that probability of an event, Λ, depends on T. Letting λ denote the annual mean arrival rate for an event,

$$\Lambda_T = 1 - e^{-\lambda T} \ .$$

We must also account for the fact that the value of the guarantee (as a fraction of consumption), θ, also depends on T. Using the annual rates g and r_f,

$$\theta_T = 1 - \left(\frac{1+r_f}{1+g}\right)^T \ .$$

With no guarantee, total expected utility at times 0 and T, i.e., $U = u(C_0) + \mathcal{E}u(C_T)(1+\delta)^{-T}$ is

$$U_1 = \frac{C_0^{1-\eta}}{1-\eta}\left\{1 + \frac{(1+g)^{(1-\eta)T}}{(1+\delta)^T}[1 - e^{-\lambda T} + (1 - e^{-\lambda T})(1-\phi)^{1-\eta}]\right\} \tag{10}$$

With the guarantee, society gives up a fraction θ_T of current and future consumption, so total utility is

$$U_2 = \left(\frac{1+r_f}{1+g}\right)^{(1-\eta)T} \frac{C_0^{1-\eta}}{1-\eta}\left\{1 + \frac{(1+g)^{(1-\eta)T}}{(1+\delta)^T}\right\} \tag{11}$$

Should society give up a fraction θ_T of consumption for the guarantee? Only if $U_2 > U_1$. U_1 and U_2 are both negative (because $\eta > 1$), so $U_2 > U_1$ implies $U_1/U_2 > 1$.

Table 1 shows the ratio U_1/U_2 for different values of η, g, λ, ϕ, and the time horizon T. It also shows θ, the value of the guarantee as a fraction of current consumption, which depends only on r_f, g, and T. Finally, it shows r_T, the effective annual discount rate for the project. (In all cases, $r_f = .01$ and $\delta = 0$.)

Note that for most of the parameter combinations shown in the table, $U_1/U_2 < 1$, so that society would not want to sacrifice the fraction θ of consumption in return for the guarantee. If the normal (i.e., ignoring a possible catastrophe) consumption growth rate is reduced from .02 to .015, the value of the guarantee falls (θ becomes .09 for $T = 20$

Table 1: Ratio of Value Functions

η	T	g	λ	ϕ	θ	r_T	U_1/U_2
2	20	0.02	0.01	0.3	.18	-.048	0.847
2	50	0.02	0.01	0.3	.39	-.014	0.634
1.5	50	0.02	0.01	0.3	.39	-.014	0.804
3	50	0.02	0.01	0.3	.39	-.014	0.392
2	20	0.015	0.01	0.3	.09	-.017	0.936
2	50	0.015	0.01	0.3	.22	-.002	0.824
2	20	0.015	0.02	0.3	.09	.013	0.960
2	50	0.015	0.02	0.3	.22	.007	0.849
2	20	0.015	0.02	0.6	.09	.048	1.097
2	50	0.015	0.02	0.6	.22	.021	1.020

Note: The guarantee should be purchased if $U_1/U_2 > 1$. The effective annual discount rate on the project is r_T, and θ is the value of the project as a fraction of current consumption. In all cases, $r_f = .01$ and $\delta = 0$. If $\lambda = .01$, $\Lambda_{20} = 0.18$ and $\Lambda_{50} = 0.39$. If $\lambda = .02$, $\Lambda_{20} = 0.33$ and $\Lambda_{50} = 0.63$.

and .22 for $T = 50$), so the ratio U_1/U_2 rises. The ratio is also higher for lower values of η. But obtaining a ratio above 1 (so that sacrificing the fraction θ of consumption is warranted) requires higher values of λ and ϕ. For example, $U_1/U_2 > 1$ if $\lambda = .02$ (so that the probability of a catastrophe occurring within 20 years is .33 and within 50 years is .63), and if the impact parameter ϕ is doubled to 0.6 (i.e., consumption drops by 60 percent if a catastrophe occurs).

Note also that the effective discount rate on the project is negative for $\lambda = .01$ and $\phi = .3$, but becomes positive if these parameters are doubled. The reason is that if a catastrophe is very likely and the impact is very large, the option to exercise the guarantee is more "in the money."

If we compare the first few rows of Table 1 with the last few, several points emerge from this extremely simple two-period example. First, buying insurance against a low-probability catastrophic event is analogous to buying an out-of-the-money put option. If that insurance is viewed as a project, it can have an effective discount rate that is negative. The value of the project (measured as a fraction of consumption, i.e., θ) can be very large, especially for long time horizons (e.g., 50 years). Furthermore, the value of the project may be larger than society's "willingness to pay" for it. But note that apart from the time horizon, the value of the project depends only on the difference between the "normal" consumption growth rate and the risk-free rate. If one believes that real consumption growth will be less than its historical average over the past 60 years (e.g., a rate of 1.5 instead of 2 percent), the value of the project is much smaller.

Also, note that to evaluate WTP, we still need values for the behavioral (or policy) parameters δ and η. As discussed, in Pindyck (forthcoming), determining the "correct" values for these parameters has been a major stumbling block for the analysis of climate change policy.

2.4. Willingness to "Build."

In Table 1, I compared the value of a project to society's willingness to pay for it. We could also compare the value with the government's willingness to "build" it. In other words, suppose that a guarantee to eliminate the impact of a catastrophe over the next 50 years is 20 percent of consumption. We could then ask whether 20 percent of consumption is sufficient revenue to enable the provider (presumably the government) to actually make (and honor) the guarantee. In the case of climate change, for example, 20 percent of consumption might or might not be sufficient to eliminate the risk of a catastrohic climate outcome.

Unlike a financial guarantee (e.g., a loan guarantee) which simply involves a possible transfer of funds, the guarantee considered here involves an actual physical investment. The cost of that investment might be so high as to make the provision of the guarantee infeasible. Or, the cost of the investment might be lower than its value, so that provision of the guarantee is a net positive NPV project.

We saw that the expected payment by the guarantor is $\phi\Lambda(1+G_T)(1-\theta)C_0 = \phi\Lambda(1+R_f)C_0$. Setting this equal to the value of the guarantee, θC_0, as in eqn. (8), gave us the effective discount rate R_T. Suppose the actual cost of making the guarantee is $K_0 = kC_0$ (i.e., expressed as a fraction of consumption). If $kC_0 > \theta C_0$, the project is economically infeasible, in the sense that its cost exceeds its value. If $kC_0 < \theta C_0$, the project will have a positive NPV, and its actual expected return will exceed R_T. Its actual expected return, R'_T, will then be:

$$R'_T = \frac{\phi\Lambda(1+R_f)}{k} - 1 \; > \; R_T \, . \tag{12}$$

Suppose $kC_0 < \theta C_0$ so the project's NPV is positive. It might be that the society's WTP for the project is less than its value θC_0 but greater than kC_0. If the project were offered by a private firm or a different government, it might then be provided at a cost just equal to the WTP, so that it still has a positive NPV. But if the government and society are one and the same, it would presumably be provided at is actual cost kC_0, so that its NPV is just zero.

In summary, we now have three numbers that characterize the economics of this project: (1) its value as a contingent claim, θC_0; (2) its actual cost, kC_0; and (3) its value to society, as measured by society's WTP. Thus evaluating the project involves a comparison of all three of these numbers.

3. Continuous Time.

Here I sketch out a simplified version of the general equilibrium model in Pindyck and Wang (forthcoming). TO BE ADDED.

References

Lucas, Deborah. 2002. "Valuing Loan Guarantees." Teaching note.

Lucas, Deborah. 2013. "Valuation of Government Policies and Projects." *Annual Review of Financial Economics.*

Pindyck, Robert S. forthcoming. "The Climate Policy Dilemma." *Review of Environmental Economics and Policy.*

Pindyck, Robert S., and Neng Wang. forthcoming. "The Economic and Policy Consequences of Catastrophes." *American Economic Journal: Economic Policy.*